Explaining Systematic Polysemy: Kinds and Individuation

Katherine Ritchie
University of California, Irvine

Sandeep Prasada
Hunter College and CUNY Graduate Center

Abstract: Polysemy is a phenomenon involving single lexical items with multiple related senses. Much theorizing about it has focused on developing linguistic accounts that are responsive to various compositional and representational challenges in semantics and psychology. We focus on an underexplored question: Why does systematic polysemy cluster in the ways it does? That is, why do we see certain regular patterns of sense multiplicity, but not others? Drawing on an independently motivated view of kind cognition—i.e., the formal structures for different classes of kind representations—we argue for an answer centered on conceptual individuation. Specifically, we argue that classes of kind concepts vary in what they individuate (i.e., counting as one and specifying what makes it the same or different from others). By elucidating these differences, we can explain why a range of patterns of systematic polysemy are found cross-linguistically and why other patterns are not attested. Overall, our account provides an explanatory framework addressing an important question at the interface between language and mind and opens new avenues for future theoretical and empirical research.

Word Count (including references and footnotes): 11,934

I. Introduction

Polysemy is a phenomenon that involves single lexical items with multiple related senses. For example, *dog* has a sense involving the kind (1a) and one involving instances of the kind (1b); *book* has an informational content sense (2a) and a physical copy sense (2b).

1. a. Dogs are widespread.

   b. Dogs are barking in the alley.

2. a. The book is depressing.

   b. The book is propping up my desk.

As we will see, these expressions fall into larger classes that display systematic and productive patterns of meaning multiplicity within and across languages.

Polysemy has long been of interest to linguists, philosophers, psychologists, and cognitive scientists. Linguists and philosophers of language have debated whether it ought to be explained as a semantic phenomena involving multiple lexicalized senses (Pustejovsky, 1995; Asher, 2011; Vicente, 2017; Devitt, 2021) or in terms of pragmatic principles like relevance or noteworthy connections in a context (Nunberg, 1995; Fauconnier, 1985; Papafragou, 1996; Wilson, 2003; Falkum, 2015; Carston, 2020). Polysemy's wide recurrence across the lexicon (and across languages), coupled with the appearance of its sense alternations tracking seemingly arbitrary properties of human conceptualization, has also been argued to undermine the possibility of an externalist referential semantics (Chomsky, 1995, 2000; Pietroski, 2018; see also, Collins, 2017).

At the interface of language and mind, the phenomenon brings to the fore questions about the nature of conceptual representation and how specifically this interacts with linguistic representation. Experimental work in psycholinguistics has been used to argue for particular mental storage profiles for polysemes involving either sense enumeration or underspecification (Klein & Murphy, 2001; Klepousniotou & Baum, 2007; Klepousniotou et al., 2012; Foraker & Murphy, 2012; MacGregor et al. 2015; Frisson, 2015; Vicente, 2018; Devitt, 2021; Löhr & Michel, 2022). Pietroski (2018) argues that polysemy is best understood as a single word providing access to multiple distinct (and potentially incompatible) concepts, one for each sense (see also Recanati, 2017); as such, the phenomenon tells us more about how language relates to conceptualization than the structure of conceptualization itself. But polysemy might also shed light on the nature of conceptual structure. For example, Quilty-Dunn (2021) argues that polysemy facts are compatible with an atomistic (rather than structured body of information) view of concepts.

Here we focus on a distinct and underexplored question: why does polysemy pattern in the ways that it does? That is, why do we find the patterns of sense alternations that we do, and not others? Our primary cases of interest display *systematic* or *regular* polysemy–a pattern in which *classes* of

expressions within and across languages display a shared sort of meaning multiplicity. Apresjan defines the notion as follows (1974: 16):

> The polysemy of a word *A* between senses a_i and a_j is *regular* in a language L if and only if there exists a word *B* in L with senses b_i and b_j, being semantically distinguished in exactly the same way as a_i and a_j, and if the pairs (a_i, b_i), (a_j, b_j) are non-synonymous. Otherwise, it is *irregular*.

For example, *dog* and *cat* both show the kind and instance-of-kind senses displayed in (1a,b), and *book* and *magazine* show the informational content and physical copy senses in (2a,b). Indeed, such patterns generalize across fairly coarse-grained semantic classes including at least biological kinds and informational artifacts. These same class-level sense alternations, and the same classes, are robustly cross-linguistically attested (Srinivasan & Rabagliati, 2015, cf. Nunberg & Zaenen, 1992).

A theory of polysemy that addresses why systematic polysemy patterns as it does must account both for the patterns we find and exclude those we don't. For example while there are regular Kind/Instance and Informational Object/Physical Object patterns of polysemy, as exemplified above, there is not a regular Vehicle/Operator pattern, evidenced by the infelicity of (3) and (4).

3. #The bus wanted the riders to sit quietly.
4. #The plane went to college in Massachusetts.

These are not felicitous means of expressing that a bus driver or pilot has the relevant properties, despite there being a systematic, cross-category, and close relationship between vehicles and their operators. We seek an account that can explain why these patterns of occurrence and nonoccurrence are productive and systematic across languages (Srinivasan & Rabagliati, 2015, 2021; Srinivasan & Snedeker, 2014).

Our question is distinct from the question of whether and how polysemy and homonymy are distinct; that question has been addressed by many linguists, psycholinguists, and philosophers (Cruse, 1986: 80; cf. Nunberg, 1979; Falkum & Vicente, 2015; Recanati, 2017; Carston 2020; Viebahn, 2018). It is also distinct from explaining why certain patterns of co-predication are (not) attested for

polysemous items. That project starts with observed patterns of polysemy, and asks how and why multiple senses of a polysemous item can or cannot be simultaneously co-predicated of a subject. For example, why is it okay to say *France is hexagonal and a republic*, but not *France is a hexagonal republic*? Proposals primarily focused on sense resolution, or on when and how senses can be combined, presuppose a space of available senses. They need not, and usually aren't, concerned with explaining why the space of systematically related senses looks the way that it does (Deane, 1988; Pustejovsky, 1995; Asher, 2011; Ortega-Andrés & Vicente, 2019; Dölling, 2020; Murphy, 2017, 2019, 2021; Löhr & Michel, 2022).

There are several accounts that get closer to our question of interest. Important attempts at developing a systematic typology that could provide the conceptual resources to encode the phenomenon have been developed (e.g., Dölling, 1995; Pustejovsky, 1995). However, specifying conceptual resources for encoding existing patterns of polysemy that are found cannot, on its own, explain why those are just the patterns that occur. Some pragmatic theorists address our question directly, relying on notions like noteworthiness in a context (e.g., Nunberg, 1995). However, such accounts are not sufficiently constrained to provide a satisfactory answer to our question, since there are plenty of noteworthy or salient relationships that do not license polysemy (Rabagliati et al., 2011).

Other theorists have argued for answers relying on conceptual representation. Srinivasan & Rabagliati (2015) argue for a partially conceptual, partially conventional account of systematic polysemy. But, they do not explain the role of concepts in constraining polysemy beyond appeal to conceptual structure. Some appeal to "co-activation" packages to explain patterns of systematic polysemy, and gesture towards relations these involve (Arapinis & Vieu, 2015; Vicente, 2015, 2017, 2021; Ortega-Andrés & Vicente, 2019; Murphy, 2021; Löhr & Michel, 2022). For example, Ortega-Andrés and Vicente offer an account in which the senses of systematic polysemes are connected by explanatory relations including "realization, actualization or implementation" (2019: 14). However,

while core explanatory relations are certainly relevant, it is not clear that they are constrained enough to explain which patterns of polysemy occur and which do not. For example, a vehicle requires an operator in order to realize its function, and yet we cannot use the term *vehicle* to refer to the operator nor vice versa. An adequate account needs to be both general enough to capture a range of relations across domains, while being constrained in ways that do not overgenerate.

We offer an answer to our question that meets this explanatory burden by appealing to the formal structure of classes of conceptual representations, rooted in psychological research and supported by resources from metaphysics. Our initial focus is on the formal characteristics of classes of kind representations which provide instructions for how to think of the contentful elements of kind representations (Prasada, 2016; Prasada & Dillingham, 2006, 2009; Haward, Carey & Prasada, 2021). We argue that these structures provide constraints that govern and explain the ways in which nouns can be polysemous between kind (1a), instance-of-kind (1b), and subkind interpretations (5).

5. The paleontologists discovered two dinosaurs—stegosauruses and velociraptors.

More broadly, we argue that these structures can provide a principled way of explaining when different forms of systematic polysemy are available. Our key insight focuses on the ways conceptual classes vary in what they *individuate*. As we use the term, individuation involves specifying as the same or different thereby allowing for something to be counted as one. Importantly, we are appealing to an epistemic or conceptual sense of individuation. This involves how we represent or think of what makes something what it is and different from others things, rather than on how something really is in the world.[1] We argue that classes of kind concepts vary in whether they individuate both instances and

[1] We do not take a stance on whether this reveals anything about the nature of objects or substances in the world. Some have focused on the metaphysics of objects or property instantiation to account for polysemy. For example, Asher takes the sorts of objects that he formally represents as having a dotted type to be metaphysically different from those that are not (see, e.g., 2011: 135). Arapinis & Vieu (2015) develop a view of polysemy on which polysemes pick out mereological composites with parts that stand in a metaphysical coincidence relation (see Gotham 2017 for a related view). See Liebesman & Magidor 2017 and 2023 for a view of the apparent polysemy of expressions like *book* that relies on the metaphysics of property instantiation. See Viebahn 2022 for challenges to Liebesman & Magidor's approach.

subkinds, whether they individuate instances in more than one way, and whether they individuate other (non-subkind) kinds. These differences, we contend, can explain a range of patterns of systematic polysemy that are found cross-linguistically, and explain why others are not.

More generally we argue that individuation can explain a wide range of patterns of systematic polysemy. We argue for the following principle:

> *(Individuation for Polysemy)* Nouns that are used to talk about *A* can also be used to talk about *B*, when a representation of *B* is formally individuated by a representation of *A*.

The crux of this principle lies in the way our notions of what something is and when it can be counted as one can depend on another sort of thing. Whenever a representation of *A* individuates a representation of *B*, we can use the term for *A* to talk about *B* as identifying one *B* intrinsically depends on what *A* is.

We begin by considering kind cognition and framing our answer to the question why systematic polysemy patterns as it does (*Sect. II*). We then build on the account, taking a closer look at representations of three informational kinds—books, magazines, and newspapers—illustrating how we can explain the variation in possible senses found among expressions for these kinds (*Sect. III*). We then consider other cases of systematic polysemy, metonymy, and deferred reference, finding the data to suggest that our proposal is on the right track (*Sect. IV*). We conclude by considering connections with other broadly conceptual views, as well as potential worries and directions for further research (*Sect. V*).

II. Kind Representations and Senses of Polysemes

We think and talk about the things we encounter in the world as instances of kinds, for example, as a dog, a table, or a tree. This seemingly simple cognitive act of thinking of something as an instance of

a kind involves the use of conceptual mechanisms whose structural or formal properties, including the centrality of principles of individuation, are often not fully appreciated. Here, we provide an overview of an independently motivated view of kind representations (Prasada, 2016; Prasada & Dillingham, 2006, 2009; Prasada et al., 2012; Rivera et al., 2023), adding some formal precision as we go.[2] The elements and mechanisms posited on this view are used for thinking and talking of things as kinds, subkinds, and instances of kinds. We argue these are precisely the sorts of conceptual structures that are needed to explain the availability of patterns of systematic polysemy. They are linked by intrinsic formal relations that enable thinking of and individuating various representations which, we will argue, license patterns of systematic polysemy.

Thinking of something as an instance of a kind using a concept like DOG not only categorizes the thing; it also grounds an explanation for why it has some of its properties and provides a way of thinking and reasoning about it from the perspective provided by its being that kind of thing (Carey, 1985; Gelman, 2003; Prasada, 2016). It leads us to think of some of an instance's properties as holding of that instance in virtue of its being the kind of thing it is (Prasada & Dillingham, 2006, 2009). For example, we understand that Fido is four-legged *because* he is a dog. The properties that are rendered intelligible by the kind are also properties that we think an instance of the kind is *supposed to* have by virtue of its being of that kind (Prasada, Khemlani, Leslie & Glucksberg, 2013). People think that Fido and other dogs are supposed to have four legs, and that if they don't, there is something wrong with them. These properties explicate what is represented as the character of the kind, distinguishing it from other kinds; call these the *characterizing properties* (CPs) associated with a kind representation.[3]

[2] We also abstract from elements of Prasada and colleagues' proposals that are not directly relevant here, but which they argue are needed to explain other aspects of how we think and talk about kinds and instances of kinds.

[3] This position does not require that concepts are represented as definitions in terms of necessary and sufficient conditions. In fact, the view of kind cognition we elaborate here has been advanced in tandem with an atomistic view of lexical concepts (Prasada & Dillingham, 2009). For discussion of definitional, atomistic, and other views of concepts, see Laurence & Margolis (1999).

CPs represent what is taken to be the non-accidental—i.e., principled—character of the kind. In many cases, we may not have a full specification of properties that characterize and distinguish a given kind from others that we represent.[4] In representing a kind, we simply assume that it has some CPs in virtue of its being the kind that it is, which differentiate it from any other kind and which we expect instances of the kind to have. That is, we assume that there is something or other that distinguishes, say, the kind elm trees from beech trees, even if we don't know what that might be (Fenton & Prasada, 2024).

Turning now to how kind representations are individuated, on the view outlined here, kind representations themselves are numerically distinct in virtue of distinct symbolic representations. Suppose someone has a kind representation K_1 and a kind representation K_2. These are distinct in virtue of their distinct symbolic representation, just as the word 'caribou' and the word 'reindeer' are distinct regardless of what they mean or are taken to pick out. When using our representations to think about the world, people do not think of distinct kinds as being merely numerically distinct and thus represented merely in terms of numerically distinct symbolic representations which allow each of them to be counted as one. Distinct kinds are understood to have distinct characters and thus kind representations have a component which represents the character of the kind via the CPs it has. These considerations suggest that the following principle of individuation is operative in kind cognition.

> *(Kind)* For any kind representations *K* and *K'*, with characterizing properties CP_K and $CP_{K'}$, a thinker counts *K* and *K'* as numerically distinct, so that each is counted as one kind, just in case:
>
> (i) $CP_K \neq CP_{K'}$ or
> (ii) when reasoning or learning about the kinds the thinker takes to be picked out by *K* and *K'* they can add a property, *P*, to CP_K without thereby adding *P* to $CP_{K'}$.[5]

[4] Putnam (1975) emphasized this point, in the service of a different aim, considering his own BEECH and ELM concepts. He did not have identifying descriptions that distinguished them, but nevertheless took them to be distinct.

[5] The second condition is needed as someone might have two kind representations with the same CPs, but nevertheless take them to pick out distinct kinds (like Putnam's (1975) description of his own BEECH and ELM concepts). The second condition allows for a way in which K and K' can be represented as being distinct, given what one expects to find out if they were to inquire further (perhaps by talking with experts). Alternatively, one might include a different condition that involved placeholders, a notion familiar from the literature on psychological essentialism (Medin & Ortony, 1989). On such a view, the CPs of kind representations would involve either distinct properties or distinct

Importantly, kind representations come in *classes*. A class of kind representations is articulated by intrinsic formal relationships (e.g., individuation) that are guaranteed by inclusion in the class. We think of these as templates that hold for concepts across the class.

Since our focus is on the formal relation of individuation, we need a clearer picture of what this involves. A principle of individuation specifies what counts as one and distinguishes it from other things of that sort.[6] If a concept, *C*, provides the formal means to represent and count things of some sort, *S*, as things of that sort, we will say *C individuates* representations of that sort *S*. For example, if a kind concept formally individuates instance-of-kind representations it enables thinking about instances as one, allowing us to count instances. Our claim is that different classes of kind representations have different formal features which individuate different sorts of representations. In the remainder of this section, we consider three classes of kind concepts–those that individuate instances and subkinds; those that individuate instances, but not subkinds; and those that individuate subkinds, but not instances. We introduce the principles of individuation for instances and subkinds in *IIA*, when introducing the class of concepts that individuates both. We'll rely on these principles further when describing the other classes of kind representations in *IIB* and *IIC*.

A. The class of kind representations that individuate instances and subkinds

A kind representation like DOG individuates—i.e., it formally enables representing as being the same or different and for counting as one—both instance representations and subkind representations.

placeholders. Placeholders would need to be distinguished modally such that placeholders *P* and *P'* are distinct if, and only if, one could "fill in" *P* without thereby filling in *P'*. We take our way of spelling out reasoning about kinds without placeholders to be more promising as a general psychological view of kind representations (even if placeholders are part of the best understanding of psychological essentialism). So, we do not appeal to placeholders in our condition on kind representation individuation, but we highlight this alternative as one that some might find promising.

[6] Our focus is not, then, on diachronic identity or tracking, say, a dog instance through changes in time. We leave open what our conceptual account suggests about such matters; see Prasada & Hall 2019 for relevant data.

Instances of DOG are formally represented as instances *of* that kind; subkinds are represented as subkinds *of* that particular kind. These formal relationships entail that instance-of-kind or subkind representations of two different kinds are distinct. That is, an instance of DOG is represented as distinct from an instance of CAT since these instances are individuated by different kind concepts.

The representations determined by a class of kind representations can differ in whether they are understood to have distinct characters (i.e., CPs) of their own, or whether they are thought to merely inherit the kind's character differing in other accidental features.

Instance-of-kind representations, like a representation FIDO, are represented as numerically distinct, but are not represented as having distinctive characters, in the way kinds are. We can represent two individual dogs that are qualitatively identical—both are small, brown, wearing a blue collar, etc.—but we cannot represent two animal kinds as having no distinguishing features (Fenton & Prasada, 2024). In our terms, this means that instance-of-kind representations for kind concepts like DOG, at least, do not have CPs. (Later, we will see classes of kind representations that distinguish other sorts of instances.) In virtue of the relationship between kind representations and instance-of-kind representations, the character of the kind is presumed to be present in all of its instances. For example, each dog instance is expected to be four-legged. The character of a kind *K* thereby specifies the manner in which we expect instances of the kind to be similar to one another (Prasada & Dillingham, 2006, 2009).

Instance-of-kind representations themselves, like kind representations, are numerically distinct in virtue of their being distinct symbolic representations. For example, the instance-of-kind representations DOG_1 and DOG_2 (or FIDO and SPOT) are distinct in terms of their symbolic representation–there are two representations of the type DOG. Since instances do not have characterizing properties, they may be distinguished in just numerical terms. However, instances are also represented as having various accidental properties (e.g., about their location, fur color, histories),

which figure in the ways people use representations to think about the world. This supports the following principle of individuation operative in physical instance-of-kind cognition.

> *(Physical Instance)* For any kind representation *K* and any physical instance representations *i* and *i'* represented with clusters of accidental properties AP_i and $AP_{i'}$, a thinker counts *i* and *i'* as numerically distinct, so that each is counted as one instance, just in case,
>
> (i) *i* is represented as an instance of *K*, *i'* is represented as an instance of *K'*, and $K \neq K'$ or
>
> (ii) $AP_i \neq AP_{i'}$ or
>
> (iii) when reasoning or learning about the individuals the thinker takes to be picked out by *i* or *i'* they can add a property *P*, to AP_i without thereby adding *P* to $AP_{i'}$.

Here, we suppose that people record clusters of accidental properties AP_i for a physical instances *i* just as they record characterizing properties CP_K for a kind *K*. Instances are distinguished (at a time) on the basis of any difference in property, including features that are not qualitative, like location or historical features.

Subkind representations, like kind representations and unlike instance-of-kind representations, are taken to have distinct characters. The characters of subkinds are taken to involve distinct ways of being their superordinate kind. The CP_S of a subkind *S* are understood to constitute a different and more specific way of realizing the character of its superordinate kind *K*. For example, COLLIE does not merely specify that its instances are a subset of the instances of DOG; all of the normative expectations we have for dogs, for example, hold of collies, but more besides. Subkinds of DOG might be represented as having a characteristic size, fur color, bark, and so on.

Subkind representations are not the same as representations of arbitrary subsets of instances. For example, while collies are understood to be one kind of dog, an arbitrary category like brown dogs are not understood to be one kind of dog. Furthermore, being a collie, more than being a brown dog, is judged to be one way of being a dog, while being a brown dog, more than being a collie, is judged to be merely some accidental feature something has in addition to being a dog (Prasada, Hennefield & Otap, 2012). COLLIE is represented as a subkind and as having a distinctive character of its own, while BROWN DOG is not.

Subkind representations are also numerically distinct in virtue of their symbolic representation. But, again, when reasoning about the world, people do not distinguish subkinds merely in terms of numerically distinct symbolic representations. Rather, they are distinguished and each is counted as one, in virtue of characterizing properties, as specified in the following principle of individuation for subkind cognition.

> *(Subkind)* For any kind representation *K* and any subkind representations *S* and *S'*, with characterizing properties CP_S and $CP_{S'}$, a thinker counts *S* and *S'* as numerically distinct, so that each is counted as one subkind just in case:
>
> (i) *S* is a represented as subkind of *K*, *S'* is represented as a subkind of *K'* and $K \neq K'$, or
>
> (ii) $CP_S \neq CP_{S'}$ or
>
> (iii) when reasoning or learning about the subkinds the thinker takes to be picked out by *S* and *S'* they can add a property *P*, to CP_S without adding thereby *P* to to $CP_{S'}$.

Kind representations in this class thereby provide three intrinsically connected forms of representation: that of the kind, that of instances of the kind, and that of subkinds of the kind. We can adopt a perspective that allows us to think and talk about the kind as a single abstract thing. Or we can adopt a perspective in which we focus on instances of a kind. Finally, we can adopt a perspective that highlights the kind's intrinsic kind-subkind structure. The formal conceptual structure connecting kind, instance-of-kind and subkind representations is retained even when one adopts, in thought or language, one or the other of those perspectives.

We propose that common nouns labeling concepts in this class, with just this formal structure, are precisely the ones that can be used to talk about kinds, instances, and subkinds. The fact that the kind representation individuates the instance and subkind representations, licenses use of the name of the kind to talk about instances and subkinds of that kind. This is because the identity and distinctness of the representations of instances and subkinds depends on the identity and distinctness of the representation of the kind. As a consequence the name of the kind can identify the kind itself and what counts as one instance and one subkind as the latter are made available by the kind representation

which formally individuates them. This formal structure correctly predicts the pattern of polysemy with nouns like *dog*, (6).

6. a. [*Kind*] Dogs evolved from wolves.

   b. [*Instance*] Two dogs that look like wolves are Fido and Rover.

   c. [*Subkind*] Two dogs that look like wolves are Huskies and German Shepherds.

Many other count nouns—especially those for animal and plant kinds—map onto kind representations with the same formal structure, and thus are also predicted to pattern like (6).

B. The class of kind representations that individuate instances, but not subkinds

A different class of kind representations have a formal structure which enables one to think of and individuate instance representations, but not subkind representations. A clear example of this class is the concept ROOK, for the piece in the game of chess. Though the kind representation formally enables representing instance-of-kind representations which are numerically distinct, it does not formally enable subkind representations that are both numerically distinct and represent principled divisions of the kind such that the subkind representations allow us to think of systematically different ways of being a rook. Because the form-function relation that characterizes rooks is an arbitrary and stipulated relation there cannot be systematically different ways of realizing the character of a rook. More generally, this class includes kinds whose instances can display only accidental qualitative differences and thus do not have potential subkinds. Other concepts within this class may include kinds that can have instances that display systematic differences making it possible to think of them as subkinds of the kind, but for which doing so is not necessary.

Consider NIGHTSTAND, for example. Suppose one asks: What is a nightstand? One reasonable sort of answer states that a nightstand is a *kind of table* that people use for some specific purposes. Such an answer reveals, at least, that we can identify the nightstands as one sort of table (or

as a way to be a table). However, nightstands can be thought of and described in ways that do not require them to be thought of as tables, as such. That is, they may not require being intrinsically thought of as one way of being a table. For example, the same question can be answered by *a nightstand is a kind of artifact that people make for*… or, *a nightstand is an artifact/something that people use for*…. And so it appears that NIGHTSTAND can be identified and individuated independently of TABLE, though, of course, it can also be explicitly identified as a kind of table

These considerations suggest that the formal structure of concepts like TABLE and ROOK do not enable one to represent subkinds: while CP_{COLLIE} specifies the particular ways in which instances realize CP_{DOG}, plausibly $CP_{NIGHTSTAND}$ merely specifies the intended functions of nightstands (see for example Heersmink, 2016).

If this is right, then the formal structure of kinds in this class should license just kind and instance readings. This is borne out in (7a-d).[7]

7. a. [*Kind*] Tables were invented by the Ancient Egyptians.

   b. [*Instance*] That store sold two tables with scratches on them.

   c. [*Subkind*] ? That store sells two tables: coffee tables and picnic tables.

   d. [*Explicitly identified subkind*] That store sells two types of tables: coffee tables and picnic tables.

The subkind interpretation is not licensed with *table* alone as it individuates instances, but not subkinds, (7c); one needs to add *type of* or *kind of* to specify the unit of quantification for that interpretation, (7d). Our view of kind representations predicts this for kinds within this class.

[7] Judgments about subkinds are subtle. We take this to be an area ripe for further empirical observation. It is also worth noting that subkind judgments differ from judgments about what we describe as abstract instances. As we will see, kinds in the informational artifactual class individuate concrete and abstract instances, but not subkinds. And while we will say that abstract instances have CPs, the informational contents of e.g. *Syntactic Structures* or *War & Peace* are not represented as 'kinds' of books. Similar things can be said about CAR: there are interesting subsets picked out by, for example, SEDAN or SUV, we hypothesize that these pattern rather like COOKBOOK with respect to BOOK; they are not represented as privileged subkinds individuated by the formal structure of CAR. Specifications of model (e.g. Camry, Corolla) pattern like *Syntactic Structures*; i.e., like abstract instances. Matters of subkind, physical instance, and abstract instance representation can be tricky here, but are relevant to asking whether and how the kind term can be used to talk about subkinds. Further systematic empirical investigation of these questions is warranted.

C. The class of kind representations that individuate subkinds, but not instances

A third class of kind representations has a structure which formally allows for thinking of and individuating subkinds, but not instances. Kind representations such as SAND and PLASTIC display this type of formal structure.

Just as with DOG, SAND specifies the characteristics that its subkinds realize in systematically different ways—in terms of color, mineral composition, particulate size, and so on. Furthermore, those subkinds cannot be thought of or characterized independently of the kind in question. Suppose one asks: what is volcanic sand? One sensible answer begins *volcanic sand is* a *kind of sand that*…. In contrast, answers beginning with *volcanic sand is a kind of stuff/substance that*… or *volcanic sand is something that*… are far less felicitous. In this way, VOLCANIC SAND patterns like COLLIE and not like NIGHTSTAND; supporting the conclusion that kind concepts in the SUBSTANCE class can individuate subkind representations.

Concepts in this class do not formally enable representing instances as being the same or different and for counting these as one. While this does not mean these concepts fail to allow us to think of portions of stuff, it does entail that concepts in this class do not provide the means to think of instances–which are countable and each quantitatively equivalent–without additional representational devices that are not specified by the formal structure of concepts like SAND or PLASTIC.[8]

[8] This does not mean that the concepts in this class, which we might think of as the SUBSTANCE class, fail to define measures for portions; cp. the interpretation of more toys and more rock (see Wellwood, forthcoming). We might sketch the relevant specification formally as follows: for K drawn from the class of kind representations that individuate subkinds but not instance, all measures μ defined for the domain of portions i of K (i.e., $D_{I(K)}$) are such that μ: DI(K) → ℝ and μ(i) = d, for some d ∈ ℝ. Many such measures will fail to support counting, including whenever the linguistic or extralinguistic context fails to specify a unit for counting. Further conditions on the selection of measures for instances in this class (e.g., Monotonicity; see ibid. for discussion and references) will ensure counting only when the stuff is appropriately contextually partitioned (cf. Sutton & Filip, 2019). See examples in 8 below and discussion there. We thank Alexis Wellwood for this suggestion and offering this way of understanding portion measurement.

Given this formal structure, we expect that nouns like *plastic* can be used to talk about the kind (8a), and its subkinds (8b), but not instances of sand (8c). In general, to talk about instances of sand requires some way of indicating what counts as one instance, e.g., through linguistic use of a measure-term like *pile*, *cup*, *pint* (8d) or through extralinguistic context (8e).

8. a. [*Kind*] Plastic was invented by chemists.

   b. [*Subkind*] The plastic used for water bottles is also used for car parts.

   c. [*Instances*] #The two plastics on the floor need to be recycled.

   d. [*Explicitly specified instances*] The two pieces of plastic on the floor need to be recycled.

   e. [*Extralinguistic specified instances*] That plastic was in my pocket. I think it was a wrapper for a sandwich.

Our proposal thereby explains, for any concept in this class, why nouns expressing those concepts can be used to talk about kinds and subkinds, but not directly, at least, about instances.

We thus see that a number of classes of basic level nouns N[9] can always be used to talk about a kind, and sometimes to talk directly about subkinds and instances. And we have said that these other levels of representation are directly accessible from N only if N expresses a kind concept *K* that provides ways of individuating instances and/or subkinds. In these cases, identifying an instance or subkind as one is dependent on the kind and the word used to identify and think about the kind can also be used to identify and think about the instances and subkinds of the kind. As such, the pattern of polysemy for a given N depends on the formal structure of *K*, which is shared by concepts within a class. This picture predicts that at least one pattern of systematic polysemy can be understood in terms of conceptual representation—the Instance/Kind form of polysemy (see Srinivasan &

[9] Nouns at this level provide the most natural answer to the question *what is that?* and tend to be at an intermediate level of classification (e.g. *dog, table, apple* rather than *animal, furniture, fruit* or *collie, dining table,* or *Macintosh apple*). Psychologists refer to it as the "basic level" as it has been shown to be psychologically privileged in a number of ways (e.g., Rosch, 1975).

Rabagliati, 2015; Asher, 2011; Vicente, 2015, 2017). We showed that classes of conceptual representation also predict the pattern involving Kind/Subkinds, rarely discussed in the polysemy literature (but see Dölling, 1995). These observations contribute to an answer to why it is that we find certain patterns of available senses amongst kind, instance and subkind senses of kind terms, and not others. We propose, however, that they also point to a significantly more general conceptual constraint on systematic polysemy.

The foregoing discussion focused on how classes of kind representations can formally individuate representations of instances or subkinds. These ideas can be carried over to explain many other patterns of systematic polysemy once we take other class-level distinctions into account, as we will see. More generally, as we stated at the outset, we propose the following principle:

> *(Individuation for Polysemy)* Nouns that are used to talk about *A* can also be used to talk about *B*, when a representation of *B* is formally individuated by a representation of *A*.

The crux of this principle lies in how representing what something is and when it can be counted as one can depend on another sort of thing. Our claim is that whenever a representation of *A* individuates a representation of *B*, this formal connection allows for the term for identifying and talking about *A* to be used to identify and talk about *B*.

An important reminder before we proceed. We have focused on how the formal notion of individuation is important for polysemy. Any of our principles might be taken to hold of non-mental things in the world—e.g., to actual dogs and tables and sand. To reiterate, however, our claim is not that patterns of systematic polysemy are determined by worldly individuation relations that hold between, e.g., particular members of an animal species and the species itself; we understand them to be determined by the ways people's *representations* of one sort figure in individuating other ones.

To illustrate our proposal further, we next consider nouns representing informational artifacts which are standardly taken to be systematically polysemous, e.g. *book* and *magazine*. Illuminating how

their formal structures are distinguished from one another, and from those concepts in the classes we have considered so far, we show how we can correctly predict the available senses of lexical items expressing these concepts.

III. A Closer Read: *Newspaper, Book,* and *Magazine*

Kind concepts like BOOK, MAGAZINE, and NEWSPAPER are taken to have instances that are material things (e.g., they can be used to prop up a table). Such instances are understood to be instances-of-*K* and as such can be taken to be qualitatively identical to an unlimited number of other instances. For example, we might represent that there are many potentially qualitatively identical copies of *Pride and Prejudice*. How these physical instances of BOOK are represented is no different than for instances of DOG or TABLE. However, BOOK, NEWSPAPER, and MAGAZINE also allow us to think of and individuate abstract instances, which are not represented as being potentially qualitatively identical. For example, a copy of *Pride and Prejudice* and a copy of *Sense and Sensibility* differ not only numerically, but also in their intelligible character, that is, in the manner in which they are understood which involves containing information, stories, or ideas, words, sentences, and images.

As we discussed above, representations themselves are numerically distinct in virtue of being represented via distinct symbolic representations. But, in reasoning about kinds, instances, and so on, other features figure in how people use their representations to think about the world. This holds for abstract instances as well. Like all instances, they are represented as distinct if they are instances of different kinds (e.g., an abstract instance representation of a magazine is represented as distinct from an abstract instance representation of a book). While physical instances may be represented as having distinct accidental properties, they are not represented as having distinct characters or CPs. Abstract instances differ in this regard. They fall into a class of representations that are distinguished by their own characterizing properties, in this case those given by what a thinker takes to be the unique

intelligible contents of a given instance. This supports that the following principle is operative in cognition about abstract instances:

> *(Abstract Instance)* For any kind representation *K* and any abstract instance representations *i* and *i'*, with characterizing properties CP_i and $CP_{i'}$, a thinker counts *i* and *i'* as numerically distinct, so that each is counted as one instance, just in case,
>
> (i) *i* is represented as an instance of *K*, *i'* is represented as an instance of *K'*, and $K \neq K'$ or
>
> (ii) $CP_i \neq CP_{i'}$ or
>
> (iii) when reasoning or learning about the instances the thinker takes to be picked out by *i* or *i'* they can add a property *P*, to CP_i without thereby adding *P* to CP_i'.

As such, concepts like BOOK, MAGAZINE, and NEWSPAPER individuate instance representations in two ways. Since instance representations are individuated in two ways, kind concepts in this class provide two ways to count instances and thus the *Individuation for Polysemy* principle predicts that the name for the kind can be used to talk about either physical or abstract instances of the kind.[10]

Kind concepts like BOOK are not, we suggest, understood to intrinsically specify subkinds. While we can identify coherent, non-arbitrary subsets of their instances on the basis of discernible differences in their physical or informational properties (e.g. cookbooks, textbooks, paperbacks, hardcover books), these are not understood to constitute merely different ways of having the CPs represented for BOOK. Rather, just as with ROOK and TABLE, we hypothesize that the properties that characterize these different subsets are taken to be explained by the different intended functions of those subtypes, rather than as subkinds that are each one way of being a rook or table. As such,

---

[10] The informational instances of BOOK, on the one hand, and those of NEWSPAPER and MAGAZINE, on the other, differ. We represent informational instances of books as fixed and unchanging, whereas informational instances of newspapers and magazines are represented as publications which themselves have distinct issues. These representational differences are borne out in differences in habitual interpretations of *book*, *newspaper,* and *magazine*.

a. Jamie reads three newspapers / magazines every week.

b. Marlena reads *Pride and Prejudice* every week.

The most natural interpretation of (a) involves reading different issues each week, while (b) can only be interpreted as saying Marlena reads the same book (maybe different copies, but the same informational content) each week.

BOOK, MAGAZINE, and NEWSPAPER do not, on their own, specify a privileged level of subkind representation or subkinds as an intrinsic unit of quantification. To talk about interesting subsets of their instances, whether physical or abstract, we must explicitly identify them via a *type of* or *kind of* locution.

Linguistic data attest to these interpretations. Applying the *Individuation for Polysemy* principle, we predict that nouns used to talk about the kinds BOOK, NEWSPAPER, and MAGAZINE, (9c), can also be used to talk about the physical instances and informational instances which these kinds figure in individuating, (9a) and (9b).

9. a. [*Instance–Physical*] The book / newspaper / magazine is tattered.

   b. [*Instance–Informational*] That book / newspaper / magazine is depressing.

   c. [*Kind*] Books were initially written on papyrus.

Since these kind representations individuate instance representations in two ways, our account predicts that they can be counted and quantified over in two ways. This is borne out in the most natural interpretations of (10)-(12).

10. a. [*Physical*] Ava recycled three magazines.

    b. [*Informational*] Ava read three magazines.

11. a. [*Physical*] Carlos shredded three old newspapers.

    b. [*Informational*] Carlos was shocked by three newspapers.

12. a. [*Physical*] Noam donated two books.

    b. [*Informational*] Noam wrote two books.

Counting data also support the prediction that representations like BOOK do not provide the formal means to individuate subkinds; compare (13) and (14)

13. [*Subkind*] ?? The two books that sell best are thrillers and romance novels.

14. [*Non-arbitrary subset*] The two types of / kinds of books that sell best are thrillers and romance novels.

Just as we saw with *table*, it is needed or at least preferred to modify *book* with a construction like *type* or *kind of* to count interesting (but not intrinsically formally privileged) subsets of book instances.

Thus far, we have considered classes of kind representations in which the noun used to talk about the kind can also be used to talk about that kind's instances and, if defined, its subkinds. Because different classes of kind representations differ in whether they formally individuate instances and subkind representations, and also in whether they individuate instance representations in more than one way, we see different patterns of polysemy across those classes.

We turn now to consider cases in which one kind representation figures in individuating a different kind and its instances. First, consider magazines and newspapers. These are produced by organizations that exist for the sake of producing the magazine or newspaper; that is, for producing the publication. These organizations are not just collections of individuals. The same individuals might be responsible for running an organization that produces one magazine at one time, and then all quit and start working for a different magazine. Or the same people might work for two different magazines. The two organizations are still represented as distinct, and counted as being two, given that they produce different products (see Noyes et al. 2023 for empirical work on distinct groups with overlapping membership). They are, we take it, represented as distinct in virtue of their distinct products—with distinct informational contents and distinct physical copies—as well as in virtue of their distinct goals to produce these two products.

Representations of organizations that produce newspapers and magazines are individuated by representations of the abstract informational product they produce. One organization is represented

as producing each product.[11] Let's return to our *Individuation for Polysemy* principle to see what is predicted. Recall that it states that nouns that are used to talk about *A* can also be used to talk about *B*, when a representation of *B* is formally individuated by a representation of *A*. Representing a publisher as being the same or different depends on representations of published product(s) being the same or different. That is to say, publisher instances are represented as individuated by the publications/products they publish. So, we predict that nouns used to talk about these products can also be used to talk about magazine and newspaper organizations or publishers. And indeed, this is what we see in (15) and (16).

15. [*Publisher Instance*] The newspaper / magazine fired its editor.
16. [*Publisher Kind*] Newspapers / Magazines have changed their fundraising efforts with the shift to online classified listings.

Magazine and *newspaper* can be used to talk about Publication-Producing Organization kinds and instances of those kinds (e.g., *The New York Times*).

Books are represented differently. They are not represented as individuating the producers (i.e., authors) who wrote them. We do not identify, distinguish or count authors by their books, and we do not suppose that producing multiple books requires multiple distinct authors. We do not distinguish an author as being many in virtue of their having authored different books.[12] Thus, on our proposal, we do not expect that the name of the product or a definite description picking it out should license talk of the producer for this class. This is borne out, (17).

17. # The book fired its editor.

[11] The single product has multiple (informational) issues each of which has multiple physical copies, but the product is still represented as a single thing.

[12] We are not claiming that people fail to accept the *Indiscernibility of Identicals* or simple consequences of it. For instance, if one discovers that *a* is *F* and *b* is *not F*, one will plausibly draw the conclusion that $a \neq b$. Our claim here is that authors are not represented as being individuated or counted as one or many in terms of their books. We are not claiming that upon discovering that my favorite author wrote a certain book that Amir's favorite author did not, people cannot thereby infer that my favorite author is not Amir's favorite author. But, this is not part of the form of the concept AUTHOR.

Book cannot be used to invoke the sense of an organization or individual who produced it.

The central role of individuation in explaining systematic polysemy is highlighted if we consider other organizations that produce products. For example, shoes are produced by companies that exist for the sake of producing shoes (perhaps in addition to other goods), but we do not represent shoes as products as individuating the companies that produce them. To further explain, let's return for a moment to the case of magazines. Magazine organizations are represented as producing a single product. Magazine organizations are represented as being distinct in virtue of producing distinct products. In contrast, shoe companies are often represented as producing many different shoe designs. Distinct shoe designs are not represented as requiring distinct shoe company producers. And, distinct shoe companies can be represented as producing the same designs—e.g., a high-top sneaker, a ballet flat, a penny loafer. Shoe instance representations (both as abstract designs and, clearly, as physical instances) do not individuate representations of shoe-producing companies as they do not specify what each is and distinguish it from others. They do not provide the means to count. Given these considerations, we do not expect that *shoe* can be used to talk about organizations that produce shoes, (18).

18. # The shoe fired its CEO.

The prediction is again borne out.

Finally, our proposal helps explain why representations of many salient and seemingly noteworthy relations fail to license polysemy. For example, magazines are often represented as having a person on their cover. This is a noteworthy, salient, and systematic relation. However, neither the representation of the magazine (as publication, informational issue, or physical copy) nor the representation of the person on the cover of an issue is represented as individuating the other. Magazines are not represented as being identified as one or distinguished as many in terms of cover people. Similarly, cover people are not represented as being identified or distinguished in virtue of

which magazines' covers they've graced. Given this, our proposal does not predict that *magazine* (or the name of a magazine) can be used to pick out a cover person, or that the name of a cover person can be used to pick out the magazine. For example, Serena Williams was on the cover of the December 2015 issue of *Sports Illustrated*. Nevertheless, *Serena Williams* cannot be used to pick out the magazine, (19a), nor can *Sports Illustrated* be used to pick out the tennis pro, (19b).

19. a. # Serena Williams is stocked next to *The Economist*.

    b. # *Sports Illustrated* played well in the tournament this weekend.

Even though there is a systematic relationship between magazines and cover people, the fact that our concepts fail to individuate either one in terms of the other explains why we should not expect a systematic Magazine/Cover Person polysemy.

The same sort of explanation applies in the case of vehicles and their operators that we considered in the introduction. Neither representations of vehicles nor of their operators are taken to individuate one another; we do not identify as being one or distinguish vehicles by their operators or operators by their vehicles; we do not count vehicles by operators or vice versa. So, while there is a systematic relationship between the two, our view provides the resources to explain why this relationship does not license a prediction of Vehicle/Operator polysemy. And, of course, it explains why terms like *sand* and *plastic* cannot be used to talk about instances of sand or plastic without the help of other concepts or nonlinguistic context to identify what counts as one instance, and also explains why *rook* cannot be used to talk about kinds of rooks.

IV. Type-M Polysemy, Context, and Metonymy

We have sketched a view of how the formal individuation structure of classes of conceptual representations can help to explain which patterns of systematic polysemy are available. Our focus has been on systematic polysemy relations involving kinds, instances, and subkinds, as well as

organizations and products. Other patterns of polysemy may have related explanations. Here, we point to several possible extensions. Before turning to further examples, though, it will be useful to draw a distinction between classes of polysemes.

Within systematic polysemy, researchers have distinguished two broad classes that we may call *Type I* and *Type M* polysemies (see, e.g., Copestake & Briscoe, 1995; Dölling, 2020). Type I polysemies are cases for which it is implausible to say that one of the associated senses is primary; their accessible meanings appear to be *inherent* (hence the use of *I*). For example, both the informational sense of *book* and its physical copy sense seem central to our conception of what books are. Similarly, the kind and instance senses of *dog* seem equally central (cf., Dölling, 2020). Type I polysemy appears to be most tightly connected to the ways we conceptualize the world, and thus constitute the cases of most interest to us.

Type M polysemies, in contrast, appear *metonymically* motivated (hence *M*) (Apresjan, 1974), and may historically have been derived from metonymic shifts (cf. Eckardt, 1999). These polysemes support the intuition that one of the senses is primary. For example, Dölling (2020) suggests that *rabbit* primarily picks out animals, with the meat or fur senses derived. Type M includes expressions that flexibly show count/mass alternations, e.g., that between Animal/Meat and Vegetable Kind/Vegetable Matter senses, as well as Container/Content alternations. The multi-sensedness of Type M polysemes arise due to regular processes of "sense extension" (see Cruse, 1986: 50f), as evidenced by, e.g., patterns of gender marking (Soler & Marti, 1993) and agreement in Spanish and Italian. Sense extension rules, for example, extend a core meaning to a new one by means of lexical (i.e., conventional) rules. And so, while these polysemes may have a (potentially partial) morphosyntactic explanation, the structure of conceptual representations could also shed light on why these relations are found systematically in human languages.

To see how this might go, consider the cases of Animal/Meat and Vegetable Kind/Vegetable Matter. Different kinds of vegetable stuff and meat are typically taken to be the same or different and counted as distinct sorts (i.e., they are individuated) by the vegetable or animal kinds from which they are derived. Potato and pumpkin matter, for example, are distinguished from one another by the kind of vegetable they are derived from; they are represented as distinct substance kinds in virtue of being derived from different vegetable kinds. Furthermore, the expected properties of the derived kind—that of the vegetable stuff or the meat—is understood to be determined by the animal or vegetable it comes from. As such, the class of ANIMAL and VEGETABLE kind representations individuate the connected substance (MEAT, VEGETABLE MATTER) kind representations. Our account would predict, then, that the names for the former can be used to talk about the latter. As such, individuation is an important component in understanding the availability of Type M polysemies.

Another case of a Producer/Product polysemy often cited in the literature, different from those we discussed in the previous section, involves artists and their creations. In this case, the roles are reversed from those we saw in the Producer/Product pattern for *newspaper* and *magazine*: here, the artist's name can be used to refer to their creative works, as in (20) and (21).

20. Two Picassos were stolen from the house.

21. Picassos are well represented in art history books.

Interestingly, this appears to only be possible when the items are understood to be creative works. One could not use (20) if what was stolen were two model airplanes that Picasso had built from a kit. We suggest that individuation helps to explain this pattern as well.[13] Artworks created by one artist are

[13] Importantly, this phenomenon is connected to, but independent of, the proper treatment of predicative occurrences of names (see Jeshion 2015 for an overview). Regardless of whether one takes argumental (*Picasso is a genius*) or predicative uses as basic (*This Picasso was born in Argentina*), the predicative uses are interpreted as properties of individuals that bear some *R* to an individual bearing the name. For arbitrary individuals, such *R*s might include 'having the same name as' or 'being relevantly similar to'. For artists and creators, the construction supports not only an *R* between individuals, but between the individual and their works. The general point, then, is that we can freely use names in predicative positions, but this regular morphosyntactic pattern does not tell us about the available interpretations for a given name.

represented as distinct from artworks created by another artist, just as DOG instances are distinguished from BEAR instances. As we saw when considering principles for representing instance individuation in *Section II* and *III*, other features matter for distinguishing two physical and abstract artworks from one another (e.g., colors, history, materials, chord progressions, lyrics…), but if artworks *a* and *a'* are represented as created by two different artists, that is sufficient for distinguishing them. The principles of individuation for artworks are taken to be intimately connected to the artists themselves—to the unique aesthetic sensibility which is displayed in their artworks. In contrast, an artist's aesthetic sensibility is not taken to be central to distinguishing or identifying the products produced by assembling model airplanes or IKEA bookshelves, as suggested by the failure of (20) to allow for a two model airplane interpretation. Further, since works are not relied upon when individuating producers (as we argued when considering books and authors), our account does not predict that products can be used to pick out their producers. Indeed, (22) cannot be used to convey that Beethoven was deaf.

22. ? The 9th symphony was deaf.

The individuation account explains these patterns.

Finally, consider another sort of case–deferred reference–as in canonical "ham sandwich" examples. These cases involve use of an expression to refer to an individual bearing some relation *R* to the referent of that expression, e.g. (23) (Nunberg 1979; see also Sag 1981, Nunberg 1995, 2004).[14]

23. The ham sandwich wants his check.

All else equal, (23) should simply read as a category mistake—who has ever known a ham sandwich to have desires? We do not antecedently represent a relation between e.g. the kind HAM SANDWICH

[14] It is common, though not universal, to differentiate polysemy from mere metonymy or "deferred reference" (although cf. Nunberg, 1995; Fauconnier, 1985; Wilson, 2003). Certainly deferred reference is not systematic or regular like the cases of polysemy we've focused on here.

and the kind PEOPLE. In contexts in which it is possible to access an instance falling under the latter kind using a name for the former kind, there is a contextually licensed one-to-one relation between the two. This suggests that, in certain situations, we may identify and distinguish (i.e., individuate) *ad hoc* kinds (e.g. people who ordered a ham sandwich) on the basis of a contextual relation holding between its instances and those of another kind (e.g. ham sandwiches). In such contexts, we can use the noun to talk about the ad hoc kind and its instances, (24a) and (24b).

24. a. Ham sandwiches rarely tip well.

    b. The ham sandwich wants a napkin.

In the right sorts of context, patrons at a restaurant are represented as being distinguished from others if they ordered distinct items (in addition to other features). In this case, individuation plays a role, but it heavily depends on context; so, we do not expect systematic polysemy between People/Food They Ordered.

The cases considered in this section suggest two further points. First, our *Individuation for Polysemy* proposal has broader applicability. While we do not claim to have shown that our proposal can explain the full range of systematic polysemy for nouns, and, while we acknowledge a role for morphosyntax and pragmatics in explanations of specific patterns, our conceptual proposal helps to explain and unify a range of standard examples of regular polysemy and also explains why many other conceptually similar cases do not license patterns of systematic polysemy. Second, our proposal helps to shed light on the ways in which Type I and Type M polysemies are similar and distinct. In these, as well as in "ham sandwich" cases, individuation is part of the explanation, and context may play a part in fixing the requisite relations. Certainly for Type I polysemies, grammatical context is less important and extralinguistic context even less so. Type M lies somewhere in between; explanations for the resolution of their senses must critically take into account at least grammatical and conceptual factors.

VI. Conclusion and Future Directions

We have offered and motivated an answer to an underexplored question: Why do we find certain regular patterns of sense multiplicity but not others? We have argued that the formal structure of certain classes of conceptual representation sheds light on this question. How kind representations enable us to think of and individuate as a matter of their structure is key to explaining why certain patterns of polysemy are and are not found.

At the outset, we noted that many thinkers have focused on what polysemy might say about the format of the representations (e.g., core meanings, sense enumeration) or the mechanisms by which they are navigated. Some posit qualia structures or co-activation packages which are central to the kind's representation (Pustejovsky, 1995; Ortega-Andres & Vicente, 2019; Vicente, 2017, 2021). Though we agree that such structures are relevant, they do not in themselves explain why certain patterns of polysemy are found while others are not. The explanatory relations of "realization, actualization or implementation" posited by Ortega-Andres & Vicente (2019) appear to reflect core aspects of our representations of various kinds, but they cannot explain why, for example, authors are required to actualize books, but we cannot use a term for a book to refer to its author. Similarly, while differences in the specification of the "agentive quale" for MAGAZINE and BOOK provide a way to represent the fact that *magazine* but not *book* can be used to talk about the producer, such theories lack an explanation for why the qualia are specified in just these ways (Pustejovsky, 1995).

Some theorists, particularly those aiming to account for polysemy in pragmatic rather than merely lexical or other semantic terms, have highlighted the relevance of conceptual relations like noteworthiness in context, salience, similarity, or co-activation patterns (Fauconnier, 1985; Nunberg, 1995; Papafragou, 1996). Some have argued that these are too unconstrained and, so, that *any* answer to the question of why polysemes pattern as they do in terms of conceptual relations will be too unconstrained. For example, Rabagliati et al. (2011) argue that characteristics of conceptual

representations like centrality and similarity fail to predict when particular senses are and are not available. They suggest instead that conventional rules are at play. Similarly, Srinivasan & Rabagliati (2015) argue that "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", and so they claim both convention and conceptual structure are important (2015: 127). While we have not argued that convention plays no role, our conceptual condition for licensing systematic polysemy is far more constrained than one relying simply on "noteworthy" relationships or co-activation patterns.

In contrast, our conceptual licensing condition explains both why certain patterns of systematic polysemy are found and others not, for a wide range of cases. And while we have only considered data from English, we expect that our *Individuation for Polysemy* principle will also help to explain why these patterns are found cross-linguistically, absent reason to think that the formal structure of classes of concepts differs across individuals or cultures (though, of course, which kind representations are lexicalized or "in use" can differ[15]). We have also left the question of which conceptual resources can explain polysemy with sense alternations across lexical categories (e.g., noun-to-verb polysemy as in *hammer*) open for future research. By appealing to the structure of kind representations in general, however, we provide an important wedge into understanding why systematic polysemy patterns as it does and open avenues for new theoretical and empirical research on the structure of and interactions between linguistic and non-linguistic representations.

Acknowledgments
We thank an audience at the New York Philosophy of Language Workshop and Washington University, St. Louis for helpful feedback. In particular, we thank Mark Fiocco, Josh Knobe, Ari

[15] For example, we can easily imagine two speakers to differ in whether they have formed a kind concept for a particular kind of substance, call it BLICK. It is furthermore attested that, in some languages, the root form expressing a given kind concept can be differentially lexically mass or count (compare *spaghetti* in English and Italian). The proper analysis of such cases is complex and not uncontested in the linguistics literature, so we leave it aside here. Thanks to Alexis Wellwood for discussion on this topic.

Koslow, and Tomasz Zyglewicz for useful discussions. Finally, we owe a special debt of gratitude to Alexis Wellwood, who was central in developing many of the ideas and arguments in this paper.

Works Cited

Apresjan, J. D. (1974). Regular polysemy. *Linguistics, 12*(142), 5–32.

Arapinis, A. and Vieu, L. (2015) A plea for complex categories in ontologies. *Applied Ontology* 10(3-4), 285-296.

Asher, N. (2011). *Lexical Meaning in Context: A Web of Words*. Cambridge University Press.

Bale, A. & Barner, B. (2009). The interpretation of functional heads: Using comparatives to explore the mass/count distinction. *Journal of Semantics*, *26(3)*, 217-252.

Carey, S. (1985). Conceptual Change in Childhood. Cambridge: MIT Press.

Carston, R. (2020). Polysemy: Pragmatics and sense conventions. *Mind & Language* 36: 108– 133. https://doi.org/10.1111/mila.12329

Chomsky, N. (1995). Language and Nature. *Mind* 104: 1-61.

Chomsky, N. (2000). New Horizons in the Study of Language and Mind. Cambridge: Cambridge University Press.

Collins, J. (2017). The copredication argument. *Inquiry*, 60(7), 675-702.

Cruse, D. A. (1986) Lexical semantics, Cambridge University Press, Cambridge, England.

Deane, P. (1988). Polysemy and cognition. *Lingua*, 75, 325-361.

Devitt, M. (2021). Semantic polysemy and psycholinguistics. *Mind & Language*, 36(1), 134–157.

Dölling, J. (2020). Systematic Polysemy. In The Wiley Blackwell Companion to Semantics (eds D. Gutzmann, L. Matthewson, C. Meier, H. Rullmann and T. Zimmermann). https://doi.org/10.1002/9781118788516.sem099

Dölling, J. (1995). Ontological Domains, Semantic Sorts and Systematic Ambiguity. *International Journal of Human-Computer Studies* 43, 785-807.

Eckardt, R.. (1999). Three ways to create metonymy. A study in locative readings of institution names. *Studi Italiani di Linguistica Teorica e Applicata* 28: 265-282.

Falkum, I. L. (2015). The how and why of polysemy: A pragmatic account. *Lingua* 157, 83–99.

Falkum, I.L. & Vicente, A. (2015). Polysemy: Current perspectives and approaches. *Lingua* 157, 1-16, https://doi.org/10.1016/j.lingua.2015.02.002.

Fauconnier, G. (1985). *Mental Spaces.* MIT Press. Cambridge, MA.

Fenton, T. & Prasada, S. (2024). Some evidence for the formal characteristics of kind and instance-of-kind representations. Poster presented at the 50th Meeting of the Society of Philosophy & Psychology, June 20-23, West Lafayette, IN.

Frisson, S. (2015). About bound and scary books: The processing of book polysemies. *Lingua*, 157, 17–35.

Foraker S, Murphy GL. (2012). Polysemy in Sentence Comprehension: Effects of Meaning Dominance. J Mem Lang. 67(4): 407-425. doi: 10.1016/j.jml.2012.07.010.

Gotham, M. (2017). Composing criteria of individuation. *Journal of Semantics*, 34, 333-371.

Haward, P., Carey, S. & Prasada, S. (2021). The Formal Structure of Kind Representations. *Cognitive Science*, *45*(10), e13040.

Heersmink, R. (2016). The metaphysics of cognitive artifacts. *Philosophical Explorations*, *19(1)*, 78-93.

Jeshion, R. (2015). Referentialism and predicativism about proper names. *Erkenntnis, 80*, 363-404.

Klein, D. E. & Murphy, G. L. (2001). The representation of polysemous words. *Journal of Memory and Language*, 45(2), 259–282.

Klepousniotou, E. & Baum, S. R. (2007). Disambiguating the ambiguity effect in word recognition: An advantage for polysemous but not homonymous words. *Journal of Neurolinguistics, 20*, 1–24.

Klepousniotou, E., Pike, G. B., Steinhauer, K. & Gracco, V. (2012). Not all ambiguous words are created equal: An EEG investigation of homonymy and polysemy. *Brain and Language*, 123, 11–21.

Laurence, S. & Margolis, E. (1999). Concepts and Cognitive Science, in E. Margolis & S. Laurence (eds.), *Concepts: Core Readings*, Cambridge, MA: MIT Press, pp. 3–81.

Liebesman, D., & Magidor, O. (2017). Copredication and Property Inheritance. *Philosophical Issues*, *27*(1), 131–166. https://doi.org/10.1111/phis.12104

Liebesman, D., & Magidor, O. (2023). Copredication and Meaning Transfer. *Journal of Semantics*, *40*, 69-91.

Löhr, G. & Michel, C. (2022). Copredication in context: A predictive processing approach. *Cognitive Science*, *46*, 1-27.

MacGregor, L., Bouwsema, J. & Klepousniotou, E. (2015). Sustained meaning activation for polysemous but not homonymous words: Evidence from EEG. *Neuropsychologia*, 68, 126–138.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In *Similarity and analogical reasoning* (pp. 179–195). Cambridge University Press. https://doi.org/10.1017/CBO9780511529863.009

Murphy, E. (2017). Acquiring the impossible: Developmental stages of copredication. *Frontiers in Psychology*, 8, 1072. https://doi.org/10.3389/fpsyg.2017.01072

Murphy, E. (2019). Acceptability properties of abstract senses in copredication. In M. Bolognesi & G. Steen (Eds.). *Perspectives on abstract concepts: From cognitive processing to semantic representation and linguistic expression* (pp. 145–165). Amsterdam: John Benjamins.

Murphy, E. (2021). Predicate order and coherence in copredication. *Inquiry*, https://doi.org/10.1080/0020174X.2021.1958054

Noyes, A., Keil, F., Dunham, Y., & Ritchie, K. (2023). *Same people, different group: Social structures are a central component of group concepts*. https://doi.org/10.31234/osf.io/a3bg2

Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3(2),143–184.

Nunberg, G. (1995). Transfers of Meaning. *Journal of Semantics*, *12*(2), 109–132. https://doi.org/10.1093/jos/12.2.109

Nunberg, G. (2004). The pragmatics of deferred interpretation. *The Handbook of Pragmatics*, 344–364.

Ortega-Andrés, M., & Vicente, A. (2019). Polysemy and co-predication. *Glossa: A Journal of General Linguistics*, *4*(1), Article 1. https://doi.org/10.5334/gjgl.564

Papafragou, A. (1996). On metonymy. *Lingua* 99:169–95.

Pietroski, P. (2018). *Conjoining Meanings: Semantics Without Truth Values*. OUP.

Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, 99, 73–112.

Prasada, S., & Dillingham, E. M. (2009). Representation of Principled Connections: A Window Onto the Formal Aspect of Common Sense Conception. *Cognitive Science*, *33*(3), 401–448. https://doi.org/10.1111/j.1551-6709.2009.01018.x

Prasada, S., Hennefield, L., & Otap, D. (2012). Conceptual and Linguistic Representations of Kinds and Classes. *Cognitive Science*, *36*(7), 1224–1250. https://doi.org/10.1111/j.1551-6709.2012.01254.x

Prasada, S., Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2013). Conceptual distinctions amongst generics. *Cognition*, *126*, 405–422. https://doi.org/10.1016/j.cognition.2012.11.010

Prasada, S. (2016). Mechanisms for thinking about kinds, instances of kinds and kinds of kinds. In *Core knowledge and conceptual change*, D. Barner & A.S. Baron (Eds.), Oxford University Press.

Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.

Putnam, H. (1975). The Meaning of "Meaning." *Minnesota Studies in the Philosophy of Science*, *7*, 131–193.

Quilty-Dunn, J. (2021). Polysemy and thought: Toward a generative theory of concepts. *Mind & Language*, *36*(1), 158–185. https://doi.org/10.1111/mila.12328

Rabagliati, H., Marcus, G. F., & Pylkkanen, L. (2011). Rules, Radical Pragmatics and Restrictions on Regular Polysemy. *Journal of Semantics*, *28*(4), 485–512. https://doi.org/10.1093/jos/ffr005

Recanati, F. (2017). Contextualism and Polysemy. *Dialectica*, *71* (3), pp.379-397. https://doi.org/10.1111/1746-8361.12179.

Rips, L. & Hespos, S.J. (2015). Divisions of the physical world: Concepts of objects and substances. *Psychological Bulletin*, *141(4)*, 786-811.

Rivera, S., Prasad, S., & Prasada, S. (2023). Are formal explanations mere placeholders or pointers? *Cognition*, *235*, 105407. https://doi.org/10.1016/j.cognition.2023.105407

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*(3), 192–233. https://doi.org/10.1037/0096-3445.104.3.192

Sag, I. 1981.Formal semantics and extralinguistic context. *Radical pragmatics*, ed.by Peter Cole, 273–94. New York: Academic Press.

Soler, C. and M.A. Marti (1993). Dealing with lexical mismatches. ACQUILEX working paper no. 2.4.

Srinivasan, M., & Snedeker, J. (2014). Polysemy and the taxonomic constraint: Children's representation of words that label multiple kinds. *Language Learning and Development*, *10*(2), 97-128.

Srinivasan, M., & Rabagliati, H. (2021). The implications of polysemy for theories of word learning. *Child Development Perspectives*, *15*(3), 148-153.

Srinivasan, M., & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, *157*, 124–152. https://doi.org/10.1016/j.lingua.2014.12.004

Sutton, P. and Filip, H. (2019). Counting constructions and coercion: Container, Portion and Measure Interpretations. In A.A. Spalek & M. Gotham (Eds.), *Approaches to Coercion and Polysemy*, *Oslo Studies in Language*, *10(2)*, 97-119.

Travis, C., 2008. *Occasion-Sensitivity: Selected Essays*. Oxford University Press, Oxford.

Vicente, A. (2015). The green leaves and the expert: Polysemy and truth-conditional variability. *Lingua*, *157*, 54-65.

Vicente, A. (2017). What words mean and express: Semantics and pragmatics of kind terms and verbs. *Journal of Pragmatics*, *117*, 231–244. https://doi.org/10.1016/j.pragma.2017.07.007

Vicente, A. (2018). Polysemy and word meaning: An account of lexical meaning for different kinds of content words. *Philosophical Studies, 175*(4), 947–968.

Vicente, A. (2021). Approaches to co-predication: Inherent polysemy and metaphysical relations. *Journal of Pragmatics*, *182*, 348-357.

Viebahn, E. (2018). Ambiguity and zeugma. *Pacific Philosophical Quarterly*, 99, 749-762.

Viebahn, E. (2022). Copredication, polysemy, and context-sensitivity. *Inquiry*, 65(8), 1066-1082.

Wellwood, A. (forthcoming). Acquiring a meaning for *more*. In G. Dupré, K. Stanton, R. Nefdt (Eds.), *Oxford Handbook of the Philosophy of Linguistics*.

Wilson, D. (2003). Relevance and lexical pragmatics. *Italian Journal of Linguistics* 15: 273–92.