

1 The Coin Toss

What are little boys made of?
What are little boys made of?
Snips and snails,
And puppy dog tails.
That's what little boys are made of.

What are little girls made of?
What are little girls made of?
Sugar and spice,
And everything nice.
That's what little girls are made of.

Making Up Meaning

Heads and Tails

One might wonder why a book about language and meaning intended for adults would open its first chapter with an epigraph of an old nursery rhyme written for children – and a sexist one at that. The point of including the rhyme is that language and how it fashions meaning, as this book hopes to demonstrate, is much like the children in the rhyme. Both are made of other smaller bits, some of which might not appear at first as sensible constituent ingredients of the whole. On further consideration, though, the aptness of the ingredients may become more apparent.¹

For the rhyme, most people would agree that children are not constructed from cayenne and mollusks. But many people might allow an at least metaphorical or loose semantic truth to the rhyme's claims about boys' and girls' characteristics (e.g., boys are often hyperactive ["puppy

¹ In the case of the rhyme, not "aptness" with respect to the claimed differences between boys and girls, but rather "aptness" in *how* the odd ingredients can invoke personality characteristics (e.g., "puppy dog tails" conjuring *liveliness* or *pep*), which could of course apply to many children regardless of their location in a gender landscape.

2 How Language Makes Meaning

dog tails”] and appreciate icky things [“snips and snails”], girls can be pleasant [“everything nice”] and can have complex personalities [“sugar and spice”]). Many other people would certainly and understandably take issue with such claims. But these protests would not likely correspond to the metaphorical or semantic meaning mechanisms per se. Rather, people would probably be upset about the claimed differential and systematic alignment of these characteristics to boys versus girls, and/or that a catchy tune aimed at children perpetuates those gender stereotypes.

With regard to language, the main idea of this book is that the way language conveys meaning is best conceptualized as a balance between two primary sets of ingredients (rather than the three supposed ingredients for a child) which are effectively both oppositional and codependent at the same time. On one side, we have a system where words and phrases correspond to other things semantically, syntactically, and semi-symbolically. The lexical item “tree,” for instance, refers to something conjured in our heads, **tree**.² This in turn corresponds to something in the external world, TREE. Such a correspondence has been discussed in various ways throughout the history of thought about language. For instance, by one simple account, spoken or written words are symbols (“tree”) for real things (TREE). Other accounts get more elaborate (see the next section, “Heads”).

This semantic/syntactic/symbolic (or hereafter: sem/syn/sym) ingredient, or portion, or perhaps half – as we’ll see – of linguistic meaning has undergone a great deal of discussion in many disciplines.³ Recent accounts basing it on embodiment have also attracted considerable attention in the form of intense theoretical debate and escalating empirical evaluation (see later in this chapter and citations in the Preface). Although these ongoing theoretical arguments and empirical studies concerning embodied simulations are fascinating in themselves, that work will *not* actually be the focus of the present book – this first ingredient of linguistic meaning, despite its laudable recent development, is not being evaluated here.

The other main set of ingredients of linguistic meaning, however, has received less attention. At best, some of its component parts have been developed in a piecemeal fashion. But they’ve rarely been discussed

² The following notation will be used to distinguish between lexical items, mentalese entries, concepts, embodied simulations, and real-world entities, respectively: “tree,” M-“tree,” **tree**, T*R*E*E, and TREE.

³ The semantics arises from lexical references as well as morphosyntax, which can contribute meaning through guiding which of several different senses of things in the external world are invoked by lexical references (e.g., a noun, “climb a tree,” a modifier, “a tree climb”).

collectively as a broader category of meaning contribution, with things in common and a consistent motivating framework.⁴ They've also not always been considered on par with, and in some ways *necessarily complementary* to, sem/syn/sym meaning. The nature of these other main ingredients, and the relationship between these two “halves” of linguistic meaning, constitute the main theme of the present work. To begin to lay the foundation of this idea, first briefly consider the sem/syn/sym side of the coin in a little more depth.

Heads

Beyond the relatively simple notion that symbolic meaning involves merely symbol and symbolized, another more detailed approach argues that a word (“tree”) corresponds to an entry in the language-of-thought or mentalese (M-“tree”). This entry also shares the meaning of “tree” but without word-like components (e.g., pronounceability, a written or signed form). The mentalese entry also in turn corresponds, *somehow*, to the concept **tree**. Such concepts have been claimed as fully or partially innate (e.g., arguably, the **solidity** of objects), or acquired through experience (e.g., learning that “hot,” **hot**, HOT things “burn,” **burn**, BURN, or that a male biological full sibling is one’s “brother,” **brother**, BROTHER). Still other concepts have been claimed to be acquired externally but with an innate predisposition *to become* acquired (e.g., the human “face,” **face**, FACE).

The operative word here, though, is “*somehow*” – *somehow* a mentalese entry must be connected to a concept. One of the shortcomings of accounts based on the notion of mentalese or something similar, however, is the lack of a clear idea about such a connection. By what means can we connect a mentalese entry, M-“tree,” with a concept, **tree**, other than they just go together through repeated but arbitrary association?

One remedy would be to just reduce the two into a single entity – somehow collapse mentalese and concepts into just the latter. But then we’re still left with the problem of, *what is a concept?*⁵ Is it just a unitary idea, a notion, a mental representation? If so, then how are concepts formed? Where do they come from? How are they distinguished and categorized? How are they manifest in the brain, etc.? We essentially have no means of *grounding* meaning easily in this kind of account without answers to these questions.

⁴ With some noted successful exceptions (e.g., Relevance Theory).

⁵ An insightful question posed to me during my Ph.D. qualifying exam (thanks, Ray), which I’ve since had the pleasure of passing along to others (you’re welcome, Kristina).

4 How Language Makes Meaning

According to a more recent approach, though, words (“tree”) correspond to one or more of a set of *embodied simulations* we can run (T*R*E*E) (Gibbs, 2005; Bergen, 2012). These simulations themselves are our sensory and motor experiences with TREES stored as generic patterns of neural activation typically in the areas of the brain responsible for sensory and motor processing. The patterns occur *authentically* when we actually interact with real TREES (e.g., when seeing TREES, climbing them, pruning them, chopping them, walking amid them, hearing them). But the patterns can also be *simulated* without peripheral sensory or motor activation as occurs when we *imagine* things like TREE. This embodied view is very promising in its ability to ground meaning in ways that are convincing biologically,⁶ evolutionarily,⁷ and neurologically⁸ – ways that ring true when contemplating the subjective (as well as shared) experience of meaning.

Grounding meaning in embodied simulations also succeeds when we consider the many decades’ worth of research in an array of areas of cognitive psychology which address embodied simulations without having necessarily sought to investigate them as such. This research aligns nicely with the idea of embodied simulations underlying certain perceptual and cognitive functioning and their characteristics (e.g., imagination, mental imagery, mental rotation, representational momentum, priming, the Perky Effect). A great deal of recent research has also determined the extent to which embodied simulations play a direct role in language processing proper (see Bergen, 2012 for a review of both the research aligning with cognitive and neuropsychology, and research demonstrating embodied simulations playing a role in language processing).

Each of these approaches – simple symbolism, mentalese, and embodied simulations – thus delineates through different means the ways in which a fixed speech signal or written form (or sign) points to, stands for, represents, or *means* somehow things in the world – often mediated by the mind doing that referring. They each thus attempt to tackle *how* those speech signals/written forms/signs achieve sem/syn/sym meaning.

But the other-side-of-meaning is much less well delineated. The relationship *between* the aforementioned sem/syn/sym side and this other-side

⁶ “Biologically convincing,” in that the fairly well-established mechanisms involved in sensation/perception and motor functioning provide a readily available set of ingredients with which to construct meaning.

⁷ “Evolutionarily convincing,” in that it makes sense that later evolving cognitive capacities (e.g., language) would have usurped already existing ones (e.g., sensation/perception, motor functioning).

⁸ “Neurologically convincing,” in that it is well-established that activation in brain regions that support *actual* sensation/perception and motor functioning occurs when those same percepts and motions are imagined, witnessed, or *contained in language being processed*.

-of-meaning is also not as thoroughly considered or discussed. Indeed, one could argue that the other-side-of-meaning and the relationship between the two sides is often miscast.

Tails

This other-side-of-meaning is in a way oppositional to the sem/syn/sym side. The sem/syn/sym side emphasizes the degree of *success* language has achieved in that words and phrases *can* mean things. The other-side-of-meaning emphasizes the degree of *failure* of this mechanism – words and phrases *don't exactly* mean things. Sem/syn/sym meaning is loaded with imprecision, ambiguity, polysemy, and other slipperiness. People don't agree universally on word/phrase meanings. If the resolution, metaphorically, on sem/syn/sym meaning is increased, the word-to-meaning correspondences become fuzzy, much like how clouds can seem whole and unitary at a distance, but foggy, edgeless, and ephemeral up close.

But, importantly, just as the success of the symbol–symbolized relationship is limited, so too is its failure of exactness. On some level, words and phrases *do* mean something. The upshot is that:

*Language works because an **optimal** level of exactness–inexactness exists in the relationship between symbol and symbolized.*

Sem/syn/sym meaning springs from the extent to which symbol *can* indicate symbolized. But this source of meaning is limited. The remaining contribution of meaning being argued for and fleshed out by this book arises from the extent to which symbol *cannot* indicate symbolized. It thus fills in where sem/syn/sym fails. The two sources of meaning are also codependent – sem/syn/sym meaning needs something to flesh out its meaning vagary, the other side requires some kernel of anchoring to guide its progression – akin to how raindrops require a mote of dust around which to form.

This other-side-of-meaning is thus enabled essentially by the optimal inexactness with which the sem/syn/sym side operates. Regarding again the sem/syn/sym side, “tree” means TREE because, depending on which of the accounts just described one invokes – I'll use embodied simulations – we've had experiences with TREE, sensorily and motorily, that are now encoded in us as relatively fixed generic patterns of neural activation. Since these patterns are recorded (i.e., trained) neurally, they can be rerun or simulated in us when we encounter the lexical item “tree,” to give us embodied simulations, T*R*E*E. So when we hear, see signed, or read “tree,” we reenact the seeing of a TREE as a T*R*E*E [seeing of], or the climbing of a TREE as a T*R*E*E [climbing of], as if we were

6 How Language Makes Meaning

actually seeing/climbing one in that moment. These reenactments thus take the form of activation in our corresponding brain regions which are for a moment semi-independent of the environment outside our skull (i.e., whether an actual tree is in our presence). The only difference between the simulation and an authentic encounter with something is that our eyes and muscles are a bit disengaged (to put it simply) in the simulation.

But again this symbol-to-symbolized correspondence is always a bit vague. It can change with time. It can trend off into several fairly distinct directions. For instance, the OBJECT TREE, as we've discussed TREE thus far, is invoked when we simulate the noun meaning of "tree" as in, "Look at that beautiful maple tree." But "tree" can also correspond to the act of PUTTING SOMETHING INTO A TREE, which we haven't discussed as much thus far, but which is entirely viable, invoked when we do a simulation from "tree" as a verb, as in "to tree a kite."

The correspondence between symbol-to-symbolized is also affected by the simulations preceding it as well as the way "tree" is embedded in its surrounding morphosyntax. The correspondence also may not be the same across individuals. For instance, people from Belize might simulate P*A*L*M. People from Alberta might simulate W*H*I*T*E S*P*R*U*C*E. People from Madagascar might simulate an enormous B*A*O*B*A*B, and an Inuit person who has never left their Arctic island home might only simulate print or screen images of trees.⁹

The following chapters present the argument that this very lack of perfect, one-to-one, symbol-to-symbolized, word-to-meaning correspondence enables another entire side of meaning to exist. Put simply, the slack in sem/syn/sym meaning provides or enables different sources of meaning themselves to arise – sources of meaning that wouldn't exist if language had a tighter symbol-to-symbolized correspondence.

Indeed, it is interesting to ponder whether the degree of accuracy in the symbol-to-symbolized connection is simply at some kind of functional plateau, in that getting all these minds together with a unitary lexical item "tree" can only get us so far. Other meaning components are then just whatever happen to be additionally leverageable by language. Or is it the case instead that the symbol-to-symbolized connection is *held* to its current level of precision by the value and utility of those other contributions to meaning enabled by the lack of perfect symbol-to-symbolized connection? So, to earn the value of those other sources of meaning, the symbol-

⁹ One can also of course take this reasoning further. A retail salesperson at a shoe store might simulate a SHOE TREE. An historian might simulate a diagram of a FAMILY TREE. A generative linguist might simulate a PHRASE STRUCTURE TREE, etc.

to-symbolized meaning precision must remain at a lesser-than-maximal level.

But either way, the two sides of meaning work in tandem to build up the resulting final overall meaning which gels, if only for a moment, in the person's mind.¹⁰ This final meaning is a complex combination of semantic content and pragmatic effects – the former comprised of embodied simulations, morphosyntax, and other semantic input, the latter resulting either from that semantic content or arising semi-independently through the multiple mechanisms that produce pragmatic effects (see Colston, 2015 for a review of how these processes “leak in” to language processing proper).

Sometimes portions of these two categories of meaning contributors are elicited under a speaker's control; other times they occur emergently. Indeed, some of the meaning that results might morph into its final form over the ongoing interaction of the interlocutors as it unfolds in time through the discourse. These portions of sem/syn/sym and pragmatic effects also combine or interact in complex ways. These braidings can also get very intricate, and can be influenced by psychological/cognitive phenomena that function in parallel to, but can also operate semi-independently from, language functioning per se. As I put it in my 2015 book, *Using Figurative Language*:

... a speaker might intend a given pragmatic effect, and that effect might not occur either because it is overtaken by some other effect or issue or because a hearer just does not compute it. But the relationship between intentionality and pragmatic meaning is much more complex than that. Pragmatic effects unintended by a speaker might occur in a hearer. A speaker might intend pragmatic effect X but instead achieve pragmatic effect Y. Hearers may be particularly primed to compute some sets of pragmatic effects but not others, resulting in a systematic bias in how intentionality is skewed – only certain effects from a speaker, intended or not, get achieved, whereas others do not occur, whether or not intended by the speaker.

The ways in which pragmatic effects can cascade off one another ... also can be intentional or not. A speaker may intend a family of pragmatic effects to arise from a figurative construction, perhaps as a logical chain sequence or spreading kind of activation. These effects then may or may not actually occur, individually or as an entirety. As with individual effects, entire sets of pragmatic effects also might arise unintentionally.

Intentionality can also emerge in the midst of a discourse rather than existing a priori in speakers. Speakers and hearers may begin a conversation without any particular intentions about pragmatic effects. They might instead just get talking about something rather innocuously, but then a pragmatic effect happens to occur. Perhaps one of the speakers unintentionally says something amusing.

¹⁰ Or, if not in tandem, then in opposition, in competition, in cahoots, in contrast, etc. The point being that both sides of meaning are present and playing their roles.

8 How Language Makes Meaning

The ensuing laughter then may instill in the speaker the more specific intention to make the other person laugh as part of the emergent camaraderie and even make the interlocutors believe that such a goal was in place at the beginning. Given the complexity of how some pragmatic effects interact with each other, a modicum of randomness also can enter into pragmatic effect computation that can assist this emergent intentionality phenomenon. (Colston, 2015, pp. 222–223)

... comprehension products, interpretative products, context, and pragmatic effects all can interact in very complex ways. They can fade and rearise and interact with schematic knowledge and potentially inaccurate memory content to influence the resulting conscious experiences of all interlocutors. Moreover, this schematic nature of memory is but one of potentially dozens of related effects concerning the malleability, alterability, inaccuracy, fallibility, and many other effects of human memory. (Colston, 2015, p. 94)

The end result of all this is an *optimal* leveraging of a maximum degree of overall final meaning being afforded at core by the individual contribution and interaction between the two basic sides of meaning.¹¹ But the contribution of the *non sem/syn/sym* side of meaning has not always been fully appreciated. To demonstrate, consider next how *sem/syn/sym* slackness has *usually* been treated.

The Disembodied Head

Symbolic/semantic inexactness or slack is taken often as a *problem* for conceptualizing language. Given the inexactness between symbol and symbolized, how do people ever align their conscious states through language? If one person says “tree,” for example, when meaning **oak**, but another hears that “tree” and comprehends **palm**, how can people ever share meaning? This “problem” has been purportedly “surmounted,” however, in a variety of ways.

Detection and Repair First, people have means by which to detect and correct for this type of inexactness. In the same ways people we know personally who might be considered “two-faced,” in that they say one thing to us but then something contradictory to someone else (e.g., that *each* of the given addressees is the speaker’s “*one and only* BFF”), are often found out, we can also detect usually when differences in interpretation have taken

¹¹ In some cursory ways, the view proposed here resembles the old bottom-up/top-down dichotomy of language processing contributions. But the current view is much more nuanced, encompassing embodied simulations, pragmatic effects, and other effects semi-independent of language functioning that leak in to affect meaning (Colston, 2015). A much greater degree of interaction among the contributing meaning components is also allowed, as is as a degree of chaotic self-emergentness to meaning that is not solely the sum of the contributing parts but is in part driven by meaning-seeking and multiply motivated human agents.

place.¹² Any given single or small set of embodied simulations usually occurs amid a much longer discourse. This lengthier interaction affords many opportunities to detect misalignments or vagaries in meaning and then hone interpretations to more closely align.

For instance, if two non-native English speakers encounter the English *lingua franca* simile, “like a walk in the park,” as a response to an inquiry about how a job interview went, they might arrive at very different meanings. The first person, who grew up perhaps in an Asian city with beautiful sculptured public parklands laced with lovely, rolling walking paths and lush tropical scenery, might consider the job interview something pleasant, leisurely, without difficulty, etc. The second person who grew up in a country where “parks” refer to enormous, undeveloped, steppe or savannah game preserves with no facilities, great distances, little available potable water, extreme heat, and many dangerous animals, might interpret the job interview as burdensome, frightening, dangerous, requiring much preparation, etc. But if these two individuals continue their discussion of the job interview, the fact that they’ve interpreted the simile differently could become apparent. They might even discover and discuss the very underpinnings outlined above as potentially producing their different initial interpretations – things like similes are often discussed when people are learning or using a *lingua franca* (Kecskes, 2007; Kecskes & Horn, 2007). So the interlocutors may be able to iron out their different interpretations on their own.

Contextual Momentum The inexactness “problem” can also be overcome through the sheer weight of the text, speech, or other context surrounding an ambiguous utterance. Consider, for example, a speaker saying, “take a taxi.” This phrase’s meaning is not at all clear when spoken in isolation. “Take” can be simulated as active or passive, and as more or less concrete. “Taxi” can also be simulated as a process, a verb, or a noun, and the latter as real or modeled. Given these alternatives, a number of different interpretations are thus available, including among others:

- 1.1a **Flag down, enter, and ride in a hired taxi to get to a destination.**
- 1.2a **Accept a toy taxi as a gift.**
- 1.3a **Physically grab a toy taxi.**
- 1.4a **Use one of a taxi company’s taxis to drive somewhere.**
- 1.5a **Consider the idea of a taxi.**
- 1.6a **Consider the idea of a toy taxi.**
- 1.7a **Consider the process of an airplane shuttling between airport gate to runway.**

¹² And, as is also the case with close friendships, we might also be particularly motivated to find such discrepancies.

10 How Language Makes Meaning

But the utterance when spoken would likely be surrounded by a context that could disambiguate its meaning, e.g.:

- 1.1b **You are in no shape to drive. *Take a taxi.***
- 1.2b **Thanks for staying late to stock the new toy shelves. I really appreciate it. I really wish I could repay you. Hey, isn't tomorrow your kid's birthday? She likes toy cars, right? Here, *take a taxi.***
- 1.3b **I'm having trouble carrying all these toy cars. I think I'm about to drop some. Could you give me a hand and grab one? No, not that one. I've got that one. *Take a taxi.***
- 1.4b **Could you please go pick up the new meters and GPS units from Randy's Supply? It's no problem, I let employees use the cabs so long as it's for company business. *Take a taxi.***
- 1.5b **There are a lot of ugly cars on the road. *Take a taxi.* It's always yellow or green or something.**
- 1.6b **There are a lot of other kinds of cars your kid could add to her collection. *Take a taxi.* She doesn't have one of those.**
- 1.7b **Being a pilot isn't *that* hard. Parts of the job are even easy. *Take a taxi,* that's no harder than driving a bus.**

The challenge this core phrase poses for interpreters – which simulations should be made and how should they be made – is thus disambiguated to a large degree by the surrounding context. Consider the last example 1.7b. Prior to reading/hearing the target phrase, the hearer has already simulated P*I*L*O*T, along perhaps with some aspects of a pilot's job, either through cascading additional simulations or the activation of schemas related to piloting.¹³ For instance, difficult parts of piloting might be simulated (e.g., T*A*K*E*O*F*F and L*A*N*D*I*N*G), as well as easier parts (e.g., C*R*U*I*S*I*N*G), etc. So the sense of “taxi” involving “an airplane shuttling from airport gate to runway,” or something close to that semantically or schematically, has probably been already partly simulated when the hearer encounters the target phrase.

Relative Similarity in Target Content Another repair to sem/syn/sym imprecision involves the relative similarity versus difference in embodied simulations made by different people. As an earlier example argued, people from different places on Earth would likely simulate different types of trees when hearing “tree” – arguably each person would simulate the tree(s) with which they're most familiar. This could lead to a misalignment between the comprehended meaning in a hearer, compared to the intended meaning of the speaker making that utterance. One could counterargue, though, that members of the category *trees*, despite

¹³ Indeed, cascading subsequent embodied simulations may be part and parcel of what schema activation in the more traditional sense actually is.