Warm (for Winter): Inferring comparison classes in communication

Michael Henry Tessler^{1,2} and Noah D. Goodman^{2,3}

¹Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

²Department of Psychology, Stanford University

³Department of Computer Science, Stanford University

Abstract

The meanings of natural language utterances depend heavily on context. Yet, what counts as context is often only implicit in conversation. The utterance *it's warm outside* signals that the temperature outside is relatively high, but the temperature could be high relative to a number of different *comparison classes*: other days of the year, other weeks, other seasons, etc. Theories of context-sensitivity in language agree that the comparison class is a crucial variable for understanding meaning, but little is known about how a listener decides upon the comparison class. Using the case study of gradable adjectives (e.g., *warm*), we extend a Bayesian model of pragmatic inference to reason flexibly about the comparison class and test its qualitative predictions in a large-scale free-production experiment. We find that human listeners infer the comparison class by reasoning about the kinds of observations that would be remarkable enough for a speaker to mention, given the speaker and listener's shared knowledge of the world. Further, we quantitatively synthesize the model and data using Bayesian data analysis, which reveals that usage frequency and a preference for basic-level categories are two main factors in comparison class inference. This work presents new data and reveals the mechanisms by which human listeners recover the relevant aspects of context when understanding language.

Keywords: comparison class; reference class; adjectives; pragmatics; Rational Speech Act; Bayesian cognitive model; Bayesian data analysis

Word count: 7851

Introduction

A 75°F (24°C) day is warm. A 50°F (10°C) day is not. That is, unless it is Winter; 50°F (10°C) could be warm for Winter. Whether or not something is *warm* depends upon what *reference class* you consider: a warm Winter day is different than a typical warm day of the year. Deciding upon a *reference class* is an open-ended problem, articulated most explicitly in the philosophy of probability when one tries to compute a probability for a singular future event (Reichenbach, 1949; Hájek, 2007): Any singular event or entity belongs to multiple categories, giving rise to multiple possible reference classes. For example, to calculate the probability that Joe Biden will be reelected President of the United States in 2024, one would ideally construct a set of similar past events and compute the frequency of successful outcomes (e.g., re-elections) relative to the total number of outcomes (number of elections); but, what “similar” past events should be considered? There is the set of all elections involving older white males at the state and federal level in Western democracies between the years 1890-2020, the set of elections of US Presidents born in Pennsylvania, the set of elections of heads-of-state who have governed during global pandemics, among countless other possibilities. Reference classes are deployed in contexts as diverse as providing a price valuation for a house (e.g., by appealing to the prices of “similar” houses) to informing the severity of sentences for defendants convicted in the legal system (e.g., by appealing to sentences of “similar” convictions; Colyvan, Regan, & Ferson, 2001; Cheng, 2009). Even in ordinary language, the reference class problem is present (e.g., what do you mean by *warm*?), and yet human speakers and listeners seem to fill in this critical aspect of context with hardly a notice.

In natural language, the problem of reference or comparison classes¹ is front and center when interpreting *relative* statements (e.g., *warm relative to what?*). The problem thus affects a wide swath of linguistic expressions, including adjectives (*warm*, *tall*; Bartsch & Vennemann, 1974; Cresswell, 1976; Klein, 1980; Kennedy & McNally, 2005), quantifiers (e.g., *many*, *a lot*; Schöller & Franke, 2017), and, in a nonobvious way, category generalizations such as generic or habitual statements (e.g., *Robins lay eggs*, *John smokes*; Tessler & Goodman, 2019). Formally, the comparison class is a free (underspecified) variable that linguistic theories assume is filled in by context (Kennedy & McNally, 2005; Bale, 2008, 2011; Solt, 2009). Psychological and/or communicative mechanisms presumably are responsible for that “filling in” process, but how those mechanisms operate to determine a comparison class remains obscure. Comparison classes thus also serve as a case study of representations at the interface of psychological and linguistic theory.

Empirically, the question of how human listeners decide upon the comparison class has received little systematic attention. Empirical work with adults and children has primarily interrogated how judgments and interpretations of relative adjectives (e.g., *dark*, *tall*) depend upon the statistical details of a pre-determined comparison class (Barner & Snedeker, 2008; Qing & Franke, 2014b; Schmidt, Goodman, Barner, & Tenenbaum, 2009; Solt & Gotzner, 2012). Indirect evidence that the comparison class is flexibly inferred comes from studies with young children, which find that a directly-modified noun phrase (e.g., *tall*

---

¹The terms *reference classes* and *comparison classes* seemingly have independent lives in independent literatures, but they are extremely related concepts. Given that our work uses language as a way to gain insight into the reference class problem, we use the terminology *comparison classes* – familiar to the linguistics literature – throughout.

pimwit, where *pimwit* is a novel category label) can be used to constrain the kinds of objects that go into the comparison class: What counts as a *tall pimwit* depends on the distribution of heights of *pimwits* and not the heights of other categories like *daxes* (i.e., *daxes* are not included in the comparison class; Barner & Snedeker, 2008). Additionally, strong linguistic and perceptual cues can provide a signal to children as young as two-and-a-half that the comparison class can change (e.g., an objectively small mitten can be *big*, if there are many tiny mittens on the table; Ebeling & Gelman, 1988, 1994). Still, a systematic investigation into how listeners flexibly adjust the comparison class is outstanding.

In what follows, we examine the problem of inferring comparison classes in communication using gradable adjectives like *warm* or *tall* as our case study. We describe how comparison classes can be incorporated into computational models of pragmatic communication. The model predicts an intuitive interaction between a listener’s world knowledge (e.g., winter days are generally cold) and the valence (or, polarity) of the adjective heard (e.g., *warm* vs. *cold*), which contrasts with the predictions of a literal Bayesian reasoner model. The pragmatic model also makes fine-grained quantitative predictions of comparison class inferences as a function of distributional knowledge about the property in the category (e.g., plausible temperatures of Winter days) and prior beliefs about the comparison class. We test these predictions in a pre-registered, large-scale free-production experiment in which participants freely generate comparison classes when interpreting relative adjectives in context. We further investigate the structure of the comparison class prior and show that it credibly contains elements that encode a preference for basic-level categories (Rosch & Mervis, 1975) and category labels that are more common in everyday speech (i.e., usage frequency).

Computational Model

Here we articulate the computations that underly comparison class inference in communication and how that inference interfaces with standard linguistic theories of adjective meaning. When interpreting an adjectival utterance u which lacks an explicit comparison class c (e.g., *he’s tall* [relative to other c]), a listener is faced with the joint inference problem of determining the comparison class c and the property value (or, degree) x described by the adjective (e.g., the height of the referent). A listener L_1 can use their knowledge of the referent k (e.g., the category membership of the referent, such as that the referent is a basketball player) to guide both their expectations about the property value $P(x | k)$ (e.g., the plausible height of a basketball player) as well as the hypothesis space of comparison classes $P(c | k)$ (described in more detail below). A Bayesian listener would infer the likely property value (e.g., height) of the referent x and the likely comparison class c by combining these prior expectations with the likelihood that a speaker $S_1(u | x, c)$ would bother to say that the adjective applies to the referent (i.e., would the speaker bother to say that the referent is tall).

$$L_1(x, c | u, k) \propto S_1(u | x, c) \cdot P(x | k) \cdot P(c | k) \quad (1)$$

Eq. 1 describes a pragmatic listener L_1 updating their beliefs about the degree x and the comparison class c by assuming that the speaker S_1 intentionally produced an adjectival

utterance u in order to accurately communicate about the degree (e.g., height). The prior distribution of comparison classes $P(c \mid k)$ is assumed to be in common ground (i.e., known to both speaker and listener) and constrained by prior knowledge about categories and the situation (e.g., the situation that the referent is a basketball player; the category knowledge that a basketball player is a person; etc.). Comparison classes can be constructed in various ways, including sets of objects in the perceptual environment (e.g., *big relative to the things around it*) or from the hypothetical functions of an object (e.g., *this shirt is big for the doll*; Ebeling & Gelman, 1994); for simplicity, we restrict our analysis to a hypothesis space of comparison classes constructed out of a taxonomic hierarchy (i.e., *conceptual comparison classes*), with the lowest level of hierarchy being the subordinate category to which the referent belongs (e.g., a basketball player; Figure 1A). The kind of property values (or, degrees) under discussion (e.g., heights) is given by the semantics of the adjectives (e.g., *tall* $\rightarrow$ height), and the hypothesis space of degrees is informed by the category membership of the referent k (e.g., plausible heights of basketball players).²

Following work in the Rational Speech Act modeling framework (Frank & Goodman, 2012; Goodman & Frank, 2016; Scontras, Tessler, & Franke, 2018), the speaker S_1 in this model is a soft-max rational agent (with degree of rationality α) who produces utterances u in order to convey information about the degree x to a listener L_0 .

$$S_1(u \mid x, c) \propto \exp(\alpha \cdot (\ln L_0(x \mid u, c) - \text{cost}(u))) \quad (3)$$

The listener $L_0(x \mid u, c)$ in the mind of the speaker $S_0(u \mid x, c)$ is assumed to know the comparison class and is simply a mapping (described below) from prior beliefs about the degree x (e.g., height) to posterior beliefs about x given an adjectival utterance that contains an explicit comparison class (e.g., *tall for a basketball player*). The speaker, as a rational actor, balances the information conveyed by the utterance, $\ln L_0(x \mid u, c)$, with the cost of producing the utterance u when deciding what to say.³ We assume the speaker has three utterances she can say: $\{\textit{tall}, \textit{short}, \textit{silence}\}$, where silence is a semantically vacuous utterance (i.e., a null action).⁴

²More generally, we can distinguish shared beliefs from private beliefs. The hypothesis space of comparison classes is constrained by the shared beliefs of the speaker and listener: A listener will not assume the speaker is using a comparison class about which the speaker is ignorant, and a listener cannot construct a comparison class from information that they do not have. By contrast, expectations about the degree value for the referent (e.g., the referent’s height) will be guided by the totality of the listener’s beliefs — both those shared with the speaker and those not shared, the listener’s private beliefs. For example, if the listener had private knowledge that the referent in question was taking a growth-stunting medication, they would use that knowledge to guide their expectations about that person’s height. If the listener knew that the speaker did not know about the growth-stunting medication, that information would not enter in the hypothesis space of comparison classes. Formally, this more general model would be given by:

$$L_1(x, c \mid u, f, g) \propto S_1(u \mid x, c) \cdot P(x \mid f, g) \cdot P(c \mid f) \quad (2)$$

where f denotes the beliefs shared between speaker and listener, and g denotes the listener’s private beliefs. Since our paradigm does not distinguish between these kinds of beliefs, we use a simplified-form of the model, where beliefs are represented by a single variable k and can be thought of as the most specific relevant category information about the referent (e.g., the referent is a basketball player).

³For simplicity, for all of our qualitative and quantitative modeling, we assume no difference in production cost for different utterances. Hence, our model reduces to simply: $S_1(u \mid x, c) \propto L_0(x \mid u, c)^\alpha$

⁴The determination of the set of alternative utterances in a Rational Speech Act model is a degree

The listener who updates their beliefs about the degree (e.g., temperature) given a vague adjectival utterance and a fixed comparison class, $L_0(x \mid u, c)$, is model of context-sensitive adjective interpretation (Lassiter & Goodman, 2013; Qing & Franke, 2014a; Lassiter & Goodman, 2017). We use a model of a literal listener which, following standard treatment in formal semantics (e.g., Kennedy, 2007), takes the literal meaning of a gradable adjective to be simply a threshold function on the degree (e.g., $\llbracket \text{tall} \rrbracket = x > \theta$) and is used to update the listener’s prior beliefs about the degree in the comparison class $P(x \mid c)$ to produce an interpretation of the adjective, contextualized by the comparison class (Figure 1B).⁵

$$L_0(x, \theta \mid u, c) \propto \delta_{\llbracket u \rrbracket(x, \theta)} \cdot P(x \mid c) \cdot P(\theta) \quad (4)$$

$$L_0(x \mid u, c) = \int_{\theta} L_0(x, \theta \mid u, c) d\theta \quad (5)$$

Equation 4 is a model of literal listener who updates their prior beliefs about the degree given a comparison class $P(x \mid c)$ via a threshold function, represented by the Kronecker delta function $\delta_{\llbracket u \rrbracket(x, \theta)}$ that returns 1 when the utterance is true (i.e., when $x > \theta$) and 0 otherwise. Lassiter and Goodman (2013, 2017) showed how the context-sensitivity of gradable adjectives can be modeled as uncertainty about the threshold $P(\theta)$ (where θ comes from a uniform prior distribution over the support of the degree prior), which we adopt here in our literal listener model.⁶ Finally, we assume the communicative goal of using an adjective like *tall* is to convey information about the height of the referent x ; thus, the speaker model S_1 (Equation 3) chooses utterances to convey the height x to the literal listener L_0 , which we calculate by marginalizing out the threshold variable θ (Equation 5).

For convenience, we consider an idealized case where there are only two comparison classes: a relatively subordinate-level category (c_{sub}) or a relatively superordinate-level category (c_{super}) (e.g., tall relative to *basketball players* or relative to *people*): $c \in \{c_{sub}, c_{super}\}$.

of freedom on the side of the modeler. We think this set is simple and reasonable, and we note that the inferences we model are invariant to other reasonable choices of alternative utterances. Most notably, including alternatives that allow the speaker to explicitly communicate the comparison class (e.g., the speaker could have said *tall for a basketball player* or *tall for a person*) do not change the key qualitative patterns we describe.

⁵Following standard semantic treatment of antonyms, the semantics of *short* are a threshold function on a distinct threshold variable: $\llbracket u_{short} \rrbracket = \text{height} < \theta_{short}$, which is also inferred by the listener model (i.e., the listener infers a threshold for both *tall* and *short*). The pragmatic inferences about the comparison class that are the focus of this paper are invariant to whether or not the antonym is included in the alternative set. The comparison class inferences are also invariant to whether or not the antonym is assigned its own unique threshold (θ_{short}). See Lassiter and Goodman (2017) and Tessler and Franke (2018) for an exploration of the implications of assigning antonyms their own thresholds vs. negating the threshold of the positive adjective.

⁶The treatment of the threshold inference at the level of the literal listener is a point of divergence from the model of Lassiter and Goodman (2013, 2017), who instead argue for “lifting” the threshold inference to the level of the pragmatic listener (i.e., the pragmatic listener infers what is likely to be an *informative threshold*, whereas the literal listener merely infers what is likely to be *true threshold*). The treatment of threshold inference at the level of the literal listener is consistent, however, with Tessler and Goodman (2019)’s model of generic language interpretation. The question of whether the threshold inference should be thought of as a pragmatic inference is an empirical question, and one that is beyond the scope of this paper. The qualitative predictions of the model about the comparison class are similar if the threshold is lifted to the pragmatic listener in such a way that the pragmatic listener jointly reasons about the threshold θ , the degree x , and the comparison class c .

Figure 1. Model overview for the example of a listener hearing a speaker describe a basketball player as *tall*. A: A hypothesis space of comparison classes is constructed over a taxonomic hierarchy. B: A comparison class is realized as a probability distribution over the relevant degree (e.g., height; shown in black). Context-specific probabilistic interpretations of the gradable adjective *tall*—given by $L_0(x | u, c)$ —are shown in red for the different comparison classes (facets). C: Listener L_1 imagines what a speaker S_1 would say given different possible heights of the referent (x-axis) assuming different comparison classes (facets); since the listener knows the referent is a basketball player, though, heights towards the upper range of the scale are *a priori* more likely (opacity). D: Marginalizing over the *a priori* plausible heights of the referent, the speaker has a preference to say *tall* over *short* if the comparison class were *other people*. E: The pragmatic listener inverts this speaker model to infer that *people* is the more likely comparison class given that the speaker said *tall*. If the referent is described as *short*, however, the listener infers the speaker meant *short for a basketball player*. Schematic prior distribution of heights for people is a unit normal distribution $\mathcal{N}(0, 1)$ and the heights of basketball players is a right-shifted normal with smaller variance $\mathcal{N}(0.5, 0.5)$.

The choice of hierarchically structured categories is convenient because subordinate category membership entails the superordinate category membership, but all that we assume in our modeling is that both categories have some non-zero prior probability of being the comparison class. Figure 1C shows the speaker production probabilities of the three possible utterances (*tall*, *short*, silence) for each value along the degree scale (e.g., each height) for each of the two different comparison classes (facets). If the comparison class is *basketball players*, the speaker will be reluctant to produce *tall* unless the height of the referent is substantially greater than that of the average basketball player (red line shifted to the right with respect to the *person* comparison class). The listener uses Bayes’ rule to “invert” this generative model of the utterance (i.e., the speaker model), inferring the implicit comparison class. In doing so, the listener deploys their prior knowledge about the height of the referent: The listener knows the referent is a basketball player and thus their height is distributed according to the distribution of heights for basketball players (Figure 1B, bottom, black; this distribution is superimposed as an opacity on Figure 1C). Averaging over the plausible heights of a basketball player, the listener reasons that a speaker who says *tall* would be more likely to do so assuming the comparison class of *people* (Figure 1D). Thus, the listener

who hears the basketball player described as *tall* tends to think the speaker meant *tall for a person*, whereas the listener who hears a jockey described as *tall* tends to think the speaker meant *tall for a jockey* (Figure 1E). Overall, the model predicts an intuitive interaction, wherein a listener infers a more superordinate comparison class when the polarity of adjective (e.g., *tall* vs. *short*) is consistent with the listener’s background knowledge about the category (e.g., the heights of basketball players vs. jockeys, respectively) than when it is inconsistent.

In our model, a pragmatic listener reasons about what comparison class a speaker would be more likely to assume. One might wonder whether such social reasoning is necessary, or whether these inferences could result from some simpler, non-social Bayesian reasoning. We examine this question by reformulating the comparison class inference formalized by the pragmatic listener (Eq. 1) in terms of a literal listener model that updates its beliefs about the world and the comparison class only via the literal meaning of the utterance and not via pragmatic reasoning (*a la* Eq. 4).

$$L_0(x, \theta, c \mid u, k) \propto \delta_{\llbracket u \rrbracket(x, \theta)} \cdot P(x \mid c) \cdot P(\theta) \cdot P(c \mid k) \quad (6)$$

Similar to the pragmatic listener model (Eq. 1), this listener knows the category membership of the referent k and can use this knowledge to constrain the hypothesis space of comparison classes via $P(c \mid k)$ (e.g., with the knowledge that the referent is a basketball player, the listener considers comparison classes that are the same as or superordinate to the class of basketball players). Unlike the pragmatic listener model (Eq. 1), however, the literal listener version of the model does not simultaneously hold different representations of the referent in mind: The pragmatic listener has both their private representation of the referent—which informs the prior distribution of the degree $P(x \mid k)$ —and imagines a speaker who acts by assuming some comparison class is in common ground— $S_1(u \mid c)$ —where c and k may or may not refer to the same class (e.g., the listener may know the referent is a basketball player – $k = \textit{basketball players}$ – but believes the speaker was assuming a *person* comparison class – $c = \textit{people}$). The literal listener version of the model has no way of separating these representations. This alternative literal listener model also predicts an interaction between background knowledge and the adjective polarity, but the directionality of these effects is exactly the opposite from the predictions of the pragmatic model (Figure 2).

Experiment

We test the model’s main prediction that the comparison class can be flexibly adjusted based on world knowledge and pragmatic principles in a large-scale, web-based free-production experiment. In the experiment, participants rephrase a speaker’s statement, which involves a scalar adjective, in a way that makes the comparison class explicit. (A smaller-scale, forced-choice version of this task was reported in Tessler, Lopez-Brau, and Goodman (2017).) To norm the world knowledge needed in our quantitative modeling, we also measure truth judgments for statements involving scalar adjectives and an explicit comparison class (described more below and described fully in SI). Sample size, exclusion criteria, regression analysis, and cognitive model analysis were preregistered for the Comparison Class Inference task: osf.io/xuc96. All data reported, analysis scripts, code for cognitive models and web experiments can be found at github.com/mhtess/comparison-class-paper.

Figure 2. Predictions of the alternative model of a literal listener that does not represent a speaker’s representation of the context as separate from their own. This listener effectively answers the question of what is more likely: a basketball player who is tall or a person who is tall?

Methods

Participants. We recruited 837 participants from Amazon’s Mechanical Turk. This number was arrived at with the goal of estimating the probability of a subordinate vs. superordinate comparison class (in a 2AFC setting) with confidence intervals no larger than 0.20 for each unique item in the data set, which conservatively produced an estimate of roughly 50 responses per item.

Participants were restricted to those with U.S. IP addresses with at least a 95% work approval rating. In addition, participants were required to pass a simple language comprehension test that we designed in order to weed out bots and other bad-faith participants (see SI). Participants who failed this check were required to exit the experiment, and so we do not have an estimate for how many participants failed this check.

Materials. The experiment involved the interpretation of gradable adjectives that describe physical dimensions (15 adjective pairs in total, e.g., *tall/short*; *warm/cold*; see Table 1 for full list). The categories (nouns) used in the experiment were a modified subset of a set of adjective-noun pairs generated by a separate group of participants ($n = 50$) in a task designed to elicit stimuli for this experiment (see SI for details). From this set of participant-generated stimuli, we curated 90 item sets, each composed of three categories at the same level of abstraction (e.g., they were all subordinate categories of the same basic-level category) that were sensible in the contexts we used in the Comparison Class Inference experiment. Two examples from each adjective pair are shown in Table 1.

Each of the 90 item sets contains 3 relatively subordinate categories, which vary according to common-sense, general expectations given by background knowledge about a degree (or, property value). For example, a basketball player, a soccer player, and a jockey vary according to common-sense expectations about their relative heights (i.e., these categories fall on the high-, medium-, and low-ends of the height scale, respectively). Furthermore, each set of 3 subordinate categories falls under the same relatively superordinate category (e.g., basketball players, soccer players, and jockeys are all *people* [the superordinate category]).⁷ The categories within each item set are then paired with either a

⁷We refer to the level of abstraction of these categories as *relatively* subordinate or superordinate, because

positive-form or negative-form adjective (e.g., tall or short). Thus, this experiment had 270 unique categories (90 sets $\times$ 3 subordinate categories of different general expectations/background knowledge) predicated in sentences by 2 adjectives each (e.g., tall and short), for a total of 540 unique items.

Adjectives (scale)	Example subordinate classes (superordinate class)
<i>big, small</i> (size)	<i>great dane, poodle, chihuahua</i> (dogs) <i>elephant, monkey, mouse</i> (animals)
<i>tall, short</i> (height)	<i>basketball player, golfer, jockey</i> (people) <i>redwood, alpine, bonsai</i> (trees)
<i>expensive, cheap</i> (price)	<i>boots, sneakers, sandals</i> (footwear) <i>steakhouse, buffet, diner</i> (restaurants)
<i>warm, cold</i> (temperature)	<i>summer, fall, winter</i> (seasons) <i>soup, salad, ice cream</i> (food)
<i>hot, cold</i> (temperature)	<i>coffee, juice, milkshake</i> (drinks) <i>sauna, shopping mall, ice rink</i> (places)
<i>heavy, light</i> (weight)	<i>wool, cotton, silk</i> (materials) <i>rock, stick, feather</i> (objects)
<i>long, short</i> (duration / length)	<i>slacks, capris, shorts</i> (pants) <i>novel, story, poem</i> (readings)
<i>loud, quiet</i> (loudness)	<i>baby, teenager, adult</i> (people) <i>auditorium, classroom, study hall</i> (rooms)
<i>noisy, quiet</i> (loudness)	<i>horn, guitar, harp</i> (instruments) <i>powerboat, sailboat, row boat</i> (boats)
<i>light, dark</i> (luminance)	<i>day, dusk, night</i> (times of day) <i>white paint, blue paint, black paint</i> (paints)
<i>fast, slow</i> (speed)	<i>runner, skier, weight lifter</i> (athletes) <i>glider, helicopter, plane</i> (aircraft)
<i>quick, slow</i> (speed)	<i>rabbit, cat, turtle</i> (pets) <i>instant pot, frying pan, crockpot</i> (cookware)
<i>strong, weak</i> (strength)	<i>hurricane, thunderstorm, rain</i> (storms) <i>lion, dog, mouse</i> (animals)
<i>hard, soft</i> (hardness)	<i>jolly rancher, chocolate, marshmallow</i> (sweets) <i>tile, wood, carpet</i> (floor materials)
<i>wide, narrow</i> (width)	<i>boulevard, street, country lane</i> (roads) <i>truck, car, golf cart</i> (vehicles)

Table 1

Thirty of the ninety sets of adjectives and categories used in the comparison class inference experiment. Categories were curated from a set of empirically elicited noun phrases from a stimulus generation task (see SI).

Procedure. The experiment began with a warm-up trial that was designed to convey the intuition behind comparison class inference. Participants were given the example of *John says: “The Empire State Building is tall”* and asked *What do you think John meant?*; participants responded by filling-in a sentence in which the comparison class was made explicit (i.e., *The Empire State Building is tall relative to other ____*). Invalid responses to this warm-up trial were used as a basis for exclusion (see SI).

Participants then completed 36 main trials. Each main trial began with a *context sentence* that introduced the referent as a member of a subordinate category and provided

the item sets may be sets of (i) subordinate-level categories within the same basic-level category (e.g., Great Danes, chihuahuas, within dogs) or (ii) basic-level categories within the same superordinate-level category (e.g., elephants, squirrels, within animals). Later, we will explicitly model the precise level of abstraction of these categories.

an appropriate, minimal context in which the adjective could be uttered (e.g., *John sees a {basketball player, golfer, jockey}*); the same context sentence was used for all 3 categories in an item set. Then, a speaker utters an adjectival utterance predicating the adjective of a pronoun used to refer to the referent (e.g., *Mary says: “They’re tall”*).⁸ The participant was asked what they thought the speaker meant (e.g., *What do you think Mary meant?*). Participants responded by freely filling in a sentence that required an explicit comparison class:

They’re tall relative to other _____.

We used the word “other” to invoke the presupposition that the referent is a member of the comparison class. Pilot testing suggesting that omitting this word invoked many direct comparisons to singular entities (e.g., *They’re tall relative to their short friend*), which were wildly heterogeneous in nature. After the main comparison class inference trials, participants completed a memory check trial where they had to select (from a checklist of 10 options) all of the adjective–noun phrase pairs (e.g., *tall – basketball player, green – tennis ball*) that they had seen on the main trials (see SI for full details).

Results

Participants were excluded if they either responded incorrectly to the warm-up trial (answering something other than *buildings, skyscrapers, structures* or the like) or answered fewer than 7 out of 10 post-test memory check questions accurately. 750 participants (89.6%) remained after these exclusion criteria, for a total of 27,000 responses. We additionally excluded nonsensical responses (443 in total, or 1.6%) and responses that resulted from the participant inferring an unintended referent for the pronoun in the target sentence (e.g., given “Alex is in a forest and hears a woodpecker. He says: *It is loud*”, the participant infers *it* refers to the forest; 267 in total, or 1.0%). We preprocessed the remaining 26317.0 responses by correcting for misspellings (462.0; 1.8%), lemmatizing (3823.0; 14.5%), and collapsing across synonyms (660.0; 2.5%; see SI for details).

Pilot testing suggested that participants primarily provided comparison class paraphrases that were identical to the subordinate noun phrase (NP) by which the referent is introduced (*subordinate-NP*, e.g., *basketball players*) or a more superordinate category (*superordinate-NP*, e.g., *people*). We took advantage of the fact that conveying a category at least as specific as the subordinate-NP generally requires including the subordinate-NP in the response (e.g., *male basketball players* is more specific than *basketball players* and includes the substring *basketball player*)⁹ to automatically categorize responses as *subordinate* if the preprocessed response contained the subordinate-NP as a substring (15721.0, or 59.7% of all responses); all other responses were coded as *superordinate* (10596, or 40.3%). Thus, “superordinate” here refers to categories potentially at different levels of abstraction (e.g., a wrestler may be strong relative to other *wrestlers, athletes, or people*; in this case, both *athletes* and *people* would be categorized as superordinate) as well as categories along

⁸We had the speaker use a pronoun so as to not provide a strong linguistic cue as to the intended comparison class. If the referent was a person, we used the singular “they” to refer to them; otherwise, we used either “it” or “they” depending on the plurality of the referent.

⁹Negation markers (e.g., “all kinds of people *except* basketball players”) are one case where this assumption would be invalid, but no responses used negation markers in this way.

different conceptual hierarchies (e.g., the price of a particular *top-shelf liquor* could be expensive relative to other *alcoholic drinks* or relative to the same kind of top-shelf liquor at *other stores*; a *pet parrot* could be loud relative to *other pets* or relative to other *birds*; etc.).

To better understand the distribution of free responses, we performed a secondary analysis where we first categorized the superordinate responses by whether or not they contained the modal superordinate-NPs for that item.¹⁰ Of the 10596 responses categorized as “superordinate”, 6645.0 (25.2% of all responses) explicitly mention the modal superordinate-NP for that item and 3951.0 (15.0% of all responses) mention neither the subordinate nor the modal superordinate-NP (“other” responses). Of the 15% of responses that mentioned neither subordinate nor superordinate responses, 77% (3034 responses) were responses that were produced by at least three participants, suggesting that these were robust alternatives to the other more salient options (i.e., subordinate or modal superordinate). These “other” responses fell into five natural categories, in roughly equal proportions: categories at an intermediate level of abstraction between subordinate and superordinate (e.g., *melons* is intermediate between *watermelons* and *fruit*; *drinks* is intermediate between *smoothies* and *food*; 28%), categories that were superordinate to the modal superordinate response (e.g., *desserts* is superordinate to *cakes* which is superordinate to *cheesecakes*; *instruments* is superordinate to *guitars* which is superordinate to *ukeleles*; 28%), categories that were superordinate along different conceptual hierarchies than the modal superordinate category (e.g., chicken at a butcher shop that was expensive relative to other *butcher shops* as opposed to relative to other *meats*; the nighttime that was dark relative to other *places* as opposed to relative to other *times of day*; 14%), rough synonyms of the modal superordinate category (e.g., *bouquets* instead of *flower bouquets*; *methods of cooking* instead of *cookware*; 21% of “other” responses), and responses that referred directly to the scale or dimension described by the adjective, which could be construed as the most superordinate category (e.g., *expensive* relative to other *prices*; 9%). In the subsequent analyses, we treat the *other* responses that mention neither the subordinate or superordinate-NP as *superordinate*. The qualitative results do not change, however, when we exclude these *other* responses altogether (see SI for a parallel set of regression results excluding the *other* responses).

Regression analysis. Our first test is to examine how subordinate vs. superordinate comparison class inferences are influenced as a function of general expectations (i.e., background knowledge) about the category of the referent and the polarity of the adjective used to describe the referent. Specifically, we test for the qualitative pattern predicted by the comparison class inference model: an interaction between the general expectations about the category and the polarity of the adjective, which we tested with a maximal Bayesian mixed-effects model.¹¹ Further, the directionality of these effects arbitrate between the

¹⁰We experimenters guessed the modal superordinate-NPs for each item (e.g., shown in Table 1), and in 89 out of 90 cases, the pre-designated superordinate-NP was the same the empirically most popular response. For the 1 item that our initial guess was incorrect (*desserts* for the set of *smoothies*, *muffins*, and *chocolate fondue*; the modal superordinate NP was in fact *foods*), we replaced our pre-designated superordinate-NP with the empirical modal superordinate-NP (*foods*).

¹¹For the regression analysis, we predict subordinate-NP paraphrases as a function of the general expectations about the subordinate category (low, medium, high; dummy coded with the medium category as the reference level), the adjective (positive vs. negative; difference coded), and their interaction; in addition, we include the maximal mixed effects structure by-item set and by-participant that mirrors this fixed effects structure. The model is $\text{subordinate\_inference} \sim \text{gen\_expectations} * \text{adjective\_polarity} +$

Figure 3. Comparison Class Inference experimental results. Proportion of paraphrases that contained the Subordinate NPs (e.g., *basketball player*) with which the referent was introduced, as a function of the general expectations (background knowledge) listeners have about the category (x-axis; e.g., gymnasts = low, basketball players = high) and the polarity of the adjective used to describe the category (e.g., *tall* = positive, *short* = negative). Bars represent overall means and error bar is a bootstrapped 95% confidence interval. Each dot represents the mean of a single item and lines connect Subordinate NPs described with different adjectives (e.g., *tall* and *short* basketball player). Dots are jittered horizontally to improve visual clarity.

pragmatic vs. literal comparison class inference models.

As predicted by the comparison class inference model, we observe a pair of two-way interactions in the directions predicted by the pragmatic, but not the literal, listener model (Figure 3): When the subordinate category was expected to be near the high-end of the scale (e.g., *basketball player*), the positive-form adjective (e.g., *tall*) led to fewer subordinate comparison classes than the negative-form adjective (e.g., *short*) in comparison to the control, middle of the scale items (e.g., *soccer player*): posterior mean beta-weight and 95% Bayesian credible interval: $\beta = -1.34 [-1.67, -1.02]$. This interaction was the result of the high-end of the scale subordinate categories showing more subordinate comparison class inferences for the negative adjective (e.g., *short* for a basketball player) than for the positive adjective (e.g., *tall* for a basketball player): $\beta = 1.41 [1.12, 1.70]$. No preference was observed for the middle-of-the-scale, control items (e.g., *tall* for a soccer player vs. *short* for a soccer player; $\beta = -0.06 [-0.24, 0.11]$). A comparable interaction was observed for categories that were expected to be near the low-end of the scale (e.g., *jockey*): Hearing the positive-form adjective led to credibly more subordinate inferences (e.g., *tall* for a jockey) than hearing a negative-form adjective (e.g., *short* for a jockey), in comparison to the middle-of-the-scale subordinate categories: $\beta = 1.27 [0.96, 1.58]$. Again, this effect was

(1 + gen_expectations * adjective_polarity | participant) + (1 + gen_expectations * adjective_polarity | item_set).

Figure 4. Comparison class inference results for 24 of the 90 items. Two examples were selected from each unique degree scale; example items were those that exhibited the greatest and smallest variability in comparison class inferences.

driven by the behavior of the low-end of the scale subordinate categories, which showed a stronger preference for the superordinate comparison classes for negative than for positive adjectives ($\beta = -1.20 [-1.51, -0.90]$).

Our results are remarkably symmetric across the low-end vs. high-end of the scale subordinate categories. We observe an overall preference for subordinate comparison classes for the control (middle-of-the-scale) subordinate categories (e.g., *soccer players*; $\beta = 0.71 [0.32, 1.10]$) and no overall differences in this preference for items at the high-end of the scale (e.g., *basketball players*; $\beta = 0.10 [-0.19, 0.41]$) or the low-end of the scale (e.g., *gymnasts*; $\beta = 0.02 [-0.26, 0.30]$). The inferences that result from adjectives that are in conflict with a listener’s general expectations of the categories (e.g., tall gymnasts vs. short basketball players) were not different between the low-end and high-end items ($\beta = -0.18 [-0.56, 0.20]$), nor were the inferences from adjectives that were consistent with general expectations about a category (e.g., short gymnasts vs. tall basketball players; $\beta = 0.02 [-0.37, 0.41]$).

Taken together, these results point to comparison class inferences that are guided by general expectations given by background knowledge about the category and the adjective that a speaker uses to describe the referent, as predicted by the pragmatic model. Human listeners infer the most likely reference class given what their knowledge of the world tells them about the observations that would be remarkable enough for a speaker to mention.

The overall preference for a subordinate level comparison class and the magnitude of the effect of the adjective varies considerably by item, however (Figure 4). We turn to the quantitative aspect of our model to triangulate the sources of this variability.

Quantitative modeling of item-variability

Modeling approach. Our model is not only useful for understanding the basic computations that underly comparison class inference. It is a tool for characterizing, by way of the gradient in the model’s predictions, the subtleties of comparison class inference. The model’s primary capacity to predict gradient in comparison class inference is as a function of background knowledge (or, *general expectations*) about properties of categories $P(x)$ (e.g., the heights of basketball players, temperatures in Winter, etc.). Beyond knowledge about properties of categories, baseline expectations about what comparison classes speakers are likely to use $P(c | k)$ can also influence the model’s predictions and be a source of variance in inferences. For the conceptual comparison classes that we focus on here, baseline expectations could vary by the level of abstraction of the categories in question (e.g., exhibiting a basic-level bias; Rosch & Mervis, 1975) as well as the usage frequency of the NPs used to describe those categories. We thus construct a family of alternative models in order to gain insight into the sources of variability in comparison class inferences; these models differ in their parameterization of the comparison class prior: (1) a Flat Prior model, which assumes the comparison class prior assigns equal probability between subordinate and superordinate classes; (2) a Basic-level Bias (or, Intercept Only) model, which assumes an a priori preference for basic-level categories; (3) a Frequency Effect (or, Slope Only) model, which assumes that NPs that are used more frequently are more preferred as comparison classes, but without a basic-level bias; (4) a Basic-level and Frequency (or, Slope and Intercept) model, which assumes both a basic-level bias and an effect of usage frequency.¹² We estimate the distributional knowledge about properties $P(x)$, the parameters of the comparison class prior (model-variant dependent), and the speaker optimality parameter α (of the speaker model in Eq. 3) using Bayesian data analysis.

Since we infer the world knowledge priors $P(x)$, our analysis technique is similar to a *descriptive Bayesian approach* (Tauber, Navarro, Perfors, & Steyvers, 2017), in which the modeler is simply interested in estimating what the prior knowledge in the model would have to be in order to account for the inference patterns in the data. Here, however, we go beyond merely accommodating the item-level variation in the data and attempt to predict the variability in the data using independent measurements of participants’ prior knowledge of categories and properties. Measuring prior knowledge in this setting is challenging, however, since direct empirical measurements require explicit estimation of relevant quantities

¹²Because our stimuli were generated by (a separate group of) participants, we do not know *a priori* if the NPs we used in the task are basic-level categories or subordinate level categories, and hence, whether the relatively more superordinate class in our stimuli should correspond to a superordinate-level category or a basic-level category. Further, a basic-level bias could operate differently for a subordinate vs. basic-level inference than for a basic vs. superordinate level inference. Thus, we endow our data-analytic model with two regression coefficient parameters corresponding to different intercept terms (i.e., a basic-level bias term for basic vs. subordinate categories, and a basic-level bias term for basic vs. superordinate categories). Further, we introduce a Bernoulli random variable z for each NP (each relatively subordinate category) to indicate whether it is a subordinate-level term or basic-level term. Further details can be found in the SI.

or probabilities (e.g., plausible temperatures of days in winter, plausible heights of basketball players, etc.; Franke et al., 2016); while many of our domains can be reasoned about intuitively (e.g., the loudness of a diesel engine vs. an electric car), most people probably do not have access to a precise representation of the underlying scale (e.g., how many decibels is the typical sound of a diesel engine?). For this reason, we take a different approach from a canonical prior elicitation task: We take advantage of the fact that our computational model for comparison class inference includes a submodel for understanding scalar adjectives with explicit comparison classes (e.g., interpreting *an electric car is quiet relative to other cars*), and we use this submodel to simultaneously predict data from a separate experiment about these similar adjectival sentences. In this *Adjective Endorsement* task, a separate group of participants ($n = 375$) provide truth judgments for adjectives with explicit comparison classes (e.g., “Imagine a day in Winter. Do you think it is cold relative to other days of the year?”; see SI for details). We then use a joint Bayesian data analytic strategy to estimate a single set of parameters for the prior knowledge (e.g., the distribution of temperatures for days in Winter) which is shared between the two tasks.

We infer the parameters for all models using a Bayesian data analytic model that shares the world knowledge parameters between the two tasks (comparison class inference and adjective endorsements) and infers the parameters for the comparison class prior (details of which depend on the model variant: Basic-level Bias, Frequency Effect, etc...) and the speaker optimality free parameters of the RSA models (SI Figure S3). We implemented the RSA and Bayesian data analysis models in the probabilistic programming language WebPPL (Goodman & Stuhlmüller, 2014) and performed inference by running 7 Markov chain Monte Carlo (MCMC) chains with 500,000 iterations each, discarding the first 250,000 for burn-in. Convergence was checked through visual inspection of the different chains to ensure similar conclusions would be drawn from each chain independently.

Model results. All model variants were able to accommodate the Adjective Endorsement data set well (Figure 5, bottom row). This result is an important sanity check; it shows that the parameters for world knowledge learned by the model are those that maximize the fit to the Adjective Endorsement data, and thus reflect this common-sense world knowledge that this task taps into. Thus, the Adjective Endorsement task works as a kind of vague, prior elicitation, where we can learn about the world-knowledge priors by asking simple natural language queries to participants. Correspondingly, the imputed world knowledge priors inferred by the Bayesian Data Analysis model reflect intuitively reasonable general expectations about the categories (e.g., the relative heights of basketball players vs. golfers vs. jockeys; Figure 6(b)).

The predictions of the different model variants come apart for the Comparison Class Inference data (Figure 5; top row). The baseline Flat Prior model predicts variability in comparison class inferences only as a function of world knowledge about the properties, assuming all comparison classes are equally likely *a priori*; this model explains roughly 14% of the variance between items. The Frequency effect model assumes NPs with higher usage frequency will more likely be used as comparison classes and is able to explain roughly 22% of the variance. The Basic-level bias model explains roughly 71% of the variance by assuming that basic-level comparison classes may be more likely *a priori*. Finally, the maximally parameterized *Basic-level and Frequency effect* – a combination of the two alternative models – gives rise to the best model predictions in terms of variance explained (77%) and mean

Figure 5. Model fits for main experiment (Comparison Class Inference; top row) and norming experiment (Adjective Endorsement; bottom row) for four models that differ in their parameterization of the comparison class prior used inside the comparison class inference RSA model. Flat prior model assumes all comparison classes are equally likely a priori. Basic-level bias model assumes that there is a preference for a basic-level comparison class. Frequency effect assumes the prior probability of a comparison class tracks the frequency of the NP in a corpus. Basic-level and frequency effect assumes that the prior probability of a comparison class is a function of a basic-level bias and frequency. Dots represents means of the human judgments (proportion of subordinate-NP responses for Comparison Class Inference experiment; proportion endorsement for norming study) and the Maximum A-Posteriori estimate of the model’s predictions. Lines represents bootstrapped 95% confidence intervals for the data and 95% Bayesian credible intervals for the models.

squared error (Table 2), suggesting that the structure of the comparison class prior reflects both a basic-level bias and an effect of the frequency of the NP, in addition to other possible factors. We confirm that the maximally parameterized model is the best explanation of the data by comparing the marginal likelihood of the data under each model to compute a Bayes Factor (BF) as a measure of formal comparison. BFs quantify how well the model predicts the data, averaging over the prior distribution over parameters; by taking the average over the model’s prior distribution over parameters, the measure explicitly takes into account model complexity because higher complexity models have wider prior distributions over parameters (Lee & Wagenmakers, 2014). The observed data is several orders of magnitude more likely under the full model than the closest competitor model, the simpler Basic-level bias only model (Table 2).

(a) Human comparison class inference data and model predictions for ten items. The log ratio frequency of subordinate NP to superordinate NP is shown in brackets next to the x-axis labels; a more negative number corresponds to a stronger prior belief in the superordinate category as the comparison class.

(b) Imputed prior distributions over degrees for ten items. Distributions were generated from the Maximum A-Posteriori parameter values inferred by conditioning on the Comparison Class Inference and Adjective Endorsement data sets. Superordinate-level category distributions (e.g., heights of people) are assumed to be Unit Normal distributions for all item sets.

Figure 6. Quantitative modeling results for ten sets of items. See SI for full data-analytic model.

To gain further insight into the comparison class inference results, we examine the maximal model’s posterior distribution over parameters. The speaker optimality parameters for each task were inferred to be values consistent with the prior literature on RSA models — means and 95% Bayesian credible intervals: $\alpha_1 = 1.45 (1.32, 1.71)$, $\alpha_2 = 5.31(4.46, 5.70)$. We also find a positive effect of frequency of the noun phrase — $\beta_1 = 0.21(0.09, 0.21)$ — indicating that more frequent noun phrases give rise to more salient comparison classes. Our maximal model infers two parameters for the basic-level bias, one which operates if the particular item is a subordinate-level category (and thus, the more superordinate category would be the basic-level; $\beta_0^0 = -0.10(-0.17, 0.10)$) and one that operates if the item is a basic-level category (and thus, the more superordinate category is the superordinate-level; $\beta_0^1 = 1.61(1.55, 1.81)$). We can use these parameters to impute prior probabilities of subordinate, basic, and superordinate comparison classes (assuming a constant effect of usage frequency). These parameters show that basic and subordinate comparison classes are roughly equally accessible while superordinate comparison classes are substantially less likely (Figure 7).

Our fully parameterized model triangulates its inferences about the comparison class via background knowledge about the property and the comparison class prior, given by

Model	r_{CC}^2	MSE_{CC}	$r_{norming}^2$	$MSE_{norming}$	log BF
Flat Prior	0.136	0.0712	0.987	0.0032	-2,569
Frequency effect	0.222	0.069	0.98	0.0058	-3,857
Basic-level bias	0.715	0.0212	0.989	0.0028	-215
Basic-level and Frequency	0.769	0.0171	0.989	0.0028	0

Table 2

Model evaluation results. Full Basic-level and Frequency model exhibits the best fit to both data sets in terms of variance explained (r^2) and mean squared error (MSE). (log) Bayes Factor are shown with respect to the full model (i.e., negative numbers indicate positive evidence for the full basic-level and frequency model).

Figure 7. Imputed distributions over the prior probabilities of comparison classes at different levels of abstraction. Basic- and subordinate-level categories comprise *a priori* likely comparison classes, while superordinate categories are less likely to serve as comparison classes.

a basic-level bias and the corpus frequency of the subordinate NP. Figure 6 shows these sources of information for ten sets of subordinate NPs. One can see that the middle-of-the-scale control items were sometimes inferred to be closer to high-end of the scale (e.g., the *parrot* is considered light in color relative to other birds) and sometimes closer to the low-end of the scale (e.g., the *elephant* is considered slow relative to other animals; Figure 6(a)). These judgments result in different imputed world knowledge priors (Figure 6(b)) which leads to different comparison class inferences (Figure 6(a), Comparison Class Inferences). For example, the *elephant* patterns more like the *sloth* than it does the *cheetah* in its comparison class inferences: when described as *fast*, it is more likely to be considered *fast* for an *elephant* than when it is described as *slow* (similar to the effect observed for *cheetah*).

The flexibility of the fully parametrized model allows it to make a wide range of quantitative predictions, which track the human data quite well. For example, comparing the *birds* items (Fig. 6(a) 2nd column) to the *meat* items (Fig. 6(a) 3rd column), we observe a dramatic difference in the overall preference for the subordinate vs. superordinate comparison class (with the *birds* items more strongly preferring the subordinate categories e.g.,

crow, parrot, seagull; and the meat items more strongly preferring the more superordinate category e.g., meat), which can be accounted for by a stronger basic-level bias for *birds* than for *meat*. The influence of corpus frequency in the comparison class prior can be seen most clearly for item sets where there are different baseline subordinate vs. superordinate preferences within the set of subordinate categories. For example, in the *vehicles* items (Fig. 6(a) right-most column), we see a strong preference for the subordinate comparison class *trucks* in the human data; at the same time, the baseline preference is stronger for the superordinate *vehicles* when talking about *smart cars* or *sedans*. The model is able to capture this pattern through the frequency component of the prior on comparison classes (the log-frequencies are shown in parentheses next to the category label in Fig. 6(a)); *trucks* is a more common label than either *smart cars* or *sedans*, which increases the model’s baseline preference for the trucks comparison class.

The data on comparison class inferences, however, sometimes exhibits even more variability and sometimes exhibits less variability than predicted by our model. For example, the model predicts variability in the inferences from the *weak* vs. *strong teenagers* vs. *adults*, where no such variability is observed empirically (similarly for the entire item set of *spices*, see Fig. S7). Comparison classes inferences concerning the strengths of different walls (tents, condos, mansions) is largely a result of different baseline preferences: *tents* is a rather good comparison class whereas *condos* and *mansions* are less strongly preferred (Fig. S7). The variability for writing tools is also quite idiosyncratic: *pens* are largely compared to other *colors* (coded as superordinate), whereas *pencils* and *chalk* are compared to other pencils and chalk respectively as well as different *scripts* (superordinate). Taken together, these results indicate that human comparison class inferences are highly flexible, in a way that reflects distributional knowledge about properties of categories, a bias for basic-level categories, and knowledge of usage frequency; still, further variance remains to be explained even in the tightly-controlled experimental data set we collected.

General Discussion

Understanding the meanings of words often, if not always, requires appreciating the context in which those words are uttered. Yet, context is almost never articulated explicitly, but left to the listener to pragmatically reconstruct. Inferring comparison classes for relative statements (e.g., scalar adjectives like *tall*) is a case study in the larger problem of pragmatic reconstruction of context and is the communicative variant of the well-known *reference class problem* in philosophy. In this paper, we propose that listeners flexibly adjust the comparison class by using their knowledge of the world to inform them about the kinds of observations that would be remarkable enough for a speaker to mention. We implement this proposal by extending an existing language understanding model to reason about comparison classes. Its predictions are empirically confirmed in an open-ended response measure using a diverse range of linguistic stimuli. The strong fit of this model suggests that comparison class inference can be viewed as cooperative pragmatic inference.

We proposed a minimal extension to an adjective-interpretation Rational Speech Act model to allow it to flexibly reason about the implicit comparison class (e.g., *tall for a person* vs. *tall for a basketball player*). We have thus addressed a particular aspect of comparison class inference: deciding among two conceptually-based comparison classes, that we pre-specified as subordinate vs. superordinate category labels. The restriction of

our model to considering only two comparison classes was a convenient simplification, but the fact that the majority of responses naturally produced in a free-production measure were either subordinate or superordinate category labels suggests that this was also a reasonable simplification. At the same time, the variability in responses was substantially greater than our idealized model could account for: Participants freely produced comparison classes that were either more abstract or less abstract than our superordinate comparison classes as well as comparison classes that fell along different conceptual hierarchies. The comparison class prior in our model could be elaborated to encode a more structured theory about how the hypothesis space of comparison classes is constructed (e.g., via a rational clustering model that considers higher-order structures in categories and properties *a la* Kemp, Shafto, & Tenenbaum, 2012).

Going beyond general conceptual knowledge, many of the “miscellaneous” comparison classes that were elicited in our task appealed to other aspects of the context that were made available by the language in the task. For example, the sentence *Robert is at a decoration shop and looks at a statue made of bronze* brings into the context different concepts that could be parametrically combined to construct comparison classes – bronze, statues, decorations, shop – and several participants appealed to these concepts in their responses (e.g., expensive relative to *items at the shop*). These observations dovetail with prior work that investigated adjective interpretation with comparison classes in young children and which highlight the importance of linguistic cues for comparison classes. For example, for a 4-year-old, the comparison class for a “tall pimwit” is *other pimwits* and does not include other non-pimwit objects (Barner & Snedeker, 2008). Even younger children can flexibly shift between qualitatively different kinds of comparison classes, given strong linguistic cues to distinguish the intended comparison class (e.g., “Is this a *big mitten*?” vs. “Is this mitten *big for the doll*?”; Ebeling & Gelman, 1994). One key future direction then is determining the cues that are available in the naturalistic environment and how adult and child listeners use those cues to infer comparison classes. The computational machinery we present in our model of comparison class inference is general and should apply equally well to these other kinds of comparison classes, once a hypothesis space of comparison classes is determined.

In addition to the novel empirical data and the computational model of comparison class inference, this paper also presents experimental and data-analytic methodological innovations. On the experimental side, we articulated an explicit generative model for our experimental stimuli, which we deployed on human participants to construct a large and diverse set of linguistic items ($n = 540$ unique stimuli). While this procedure was not entirely “end-to-end” (i.e., we authors still needed to curate, edit, and add context to the items), the method presents a significant advance beyond the traditional method of constructing a small set of stimuli, often inadvertently optimized to test a theory (Clark, 1973). On the data-analytic side, we coupled a *descriptive Bayesian* approach (Tauber et al., 2017) with the productivity of probabilistic models of language understanding (Goodman & Frank, 2016; Scontras et al., 2018) to jointly model two complementary language tasks and infer the relevant prior knowledge shared between the tasks. The major feature of this method is that it allows us to back-out quantitatively detailed domain knowledge that would otherwise be inaccessible through traditional prior elicitation techniques because human participants lack requisite knowledge of the quantitative scales (e.g., how many decibels is the sound of a rooster’s crow?); in addition, this method has the feature that participants respond

only to simple, natural language questions rather than estimate numerical quantities for which complicated linking functions must be designed (cf., Franke et al., 2016). This fully Bayesian language approach provides further constraints on the language understanding models, which must predict quantitative data from two related language experiments. This approach highlights how the productivity of natural language can be harnessed to productively design experiments to further constrain and test computational models of language and cognition.

In our stimuli, a speaker utters an adjectival sentence without a strong cue to the comparison class: *He [a basketball player] is tall*. Why might speakers not bother articulating the comparison class? If a listener is trying to infer the comparison class (like the model we present in this paper), then a speaker (who is reasoning about this listener) could save the effort of producing the comparison class when they believe the listener will correctly infer it.¹³ In the SI, we formalize this kind of higher-order speaker, who reasons about whether or not the listener will correctly infer the comparison class; furthermore, we show that a listener who reasons about this speaker (i.e., a higher-order listener) draws qualitatively similar comparison class inferences to the simpler listener model we present above. Future work can be done to interrogate the predictions of the higher-order speaker to better understand when human speakers decide when to make the comparison class explicit.

Comparison classes are useful for learning linguistic meanings because they allow us to generalize the lexical semantics of many different kinds of adjectives and other linguistic messages that convey relative meanings. For example, the semantic meaning of *tall* and *warm* have equivalent semantic functional forms, which vary only in the dimension the adjective picks up: $\llbracket tall \rrbracket = \text{height}(x) > \theta$ and $\llbracket warm \rrbracket = \text{temperature}(x) > \theta$. Even generic language (e.g., *Birds have hollow bones*) can be understood by a threshold semantics: $\llbracket Birds \text{ fly south in the winter} \rrbracket = P(x \text{ flies south in the winter} \mid x \text{ is a bird}) > \theta$ (Tessler & Goodman, 2019). Human cognition can use the minimal semantic template – $\llbracket words \rrbracket = \text{dimension}(x) > \theta$ – and infer the intended comparison class (which determines θ) to create an infinity of possible meanings, and natural language seems to embrace these kinds of ambiguities (Piantadosi, Tily, & Gibson, 2012). It remains an open question, however, whether learning this common structural form of meaning requires a model with built-in structure. For example, large-scale deep learning (especially transformer-based) sequence models that are trained from huge amounts of text (i.e., Large Language Models, such as BERT or GPT-3; Devlin, Chang, Lee, & Toutanova, 2018, Brown et al., 2020) exhibit surprisingly sophisticated syntactic knowledge, even on syntactic constructions that are rare in natural text (Linzen, Dupoux, & Goldberg, 2016; Futrell, Wilcox, Morita, & Levy, 2018; Wilcox, Futrell, & Levy, 2021). An important future direction for this line of work is to test whether such models learn a shared semantic representation for scalar adjectives (or, relative statements more broadly) and whether the Large Language Models appreciate the context-sensitivity that comes with different comparison classes.

The phenomenon of comparison class inference is a case study in *context inference*, or reasoning about the set of beliefs that are in common ground between speaker and listener. The model is thus similar in structure to models of presupposition accommodation (e.g., Qing, Goodman, & Lassiter, 2016; Degen, Tessler, & Goodman, 2015), which linguistic

¹³The speaker here would be operating roughly according to a principle of trying to communicate more information with less effort.

theories classically pose as involving a listener adding or revising information to the common ground in order to make sense of an utterance. For example, if John says to Mary “My car is in the shop”, it is not necessary *a priori* that Mary know that John has a car; if Mary did not know (or she temporarily forgot) that John owns a car, Mary can *accommodate* John’s utterance by adding to the common ground the presupposition that *John has a car*. The continuity in the computational formalisms that underly this sort of common ground revision suggest a very basic cognitive mechanism behind the pragmatic reconstruction of context (e.g., Levinson, 2000).

Comparison class inference is also related to the well-known problem of *reference classes* (Reichenbach, 1949; Hájek, 2007). In communication, pragmatic mechanisms are at play in how listeners decide upon an appropriate comparison class. This investigation thus raises the question of whether or not similar social reasoning occurs in legal or actuarial contexts, where reference class inference can impact everything from the value of a house to the severity of a criminal sentence (Colyvan et al., 2001; Cheng, 2009). If so, the very way in which we appreciate justice being served or the value of property stems from a cooperative, communicative framework.

Conclusion

The words we say are often too vague to have a single, precise meaning; they only make sense in context. Aspects of the context, however, can also be underspecified, leaving the listener in the dark about both the speaker’s intended meaning and about the context through which the listener is to make sense of the conversation. This work suggests that listeners infer a lot from a little: Meaning and context from simple, vague utterances.

Acknowledgements

We thank Polina Tsvilodub for her contributions to the preprocessing of the free production data and for implementing the adjective endorsement task. We also thank Michael Lopez-Brau for work on a preliminary version of this project. This work was supported in part by NSF Graduate Research Fellowship DGE-114747 and NSF SBE Postdoctoral Research Fellowship Grant No. 1911790 to MHT and a Sloan Research Fellowship, ONR grant N00014-13-1-0788, and DARPA grant FA8750-14-2-0009 to NDG.

References

- Bale, A. C. (2008). A universal scale of comparison. *Linguistics and Philosophy*, 31(1), 1.
- Bale, A. C. (2011). Scales and comparison classes. *Natural Language Semantics*, 19, 169–190.
- Barner, D., & Snedeker, J. (2008). Compositionality and Statistics in Adjective Acquisition: 4-year-olds Interpret Tall and Short Based on the Size Distributions of Novel Noun Referents. *Child Development*, 79, 594–608.
- Bartsch, R., & Vennemann, T. (1974). *Semantic structures*.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... others (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Cheng, E. K. (2009). A practical solution to the reference class problem. *Colum. L. Rev.*, 109, 2081.
- Clark, H. H. (1973). The Language-as-Fixed-Effect Fallacy: A Critique of Language Statistics in Psychological Research. *Journal of Verbal Learning and Verbal Behavior*, 12, 335–359.
- Colyvan, M., Regan, H. M., & Ferson, S. (2001). Is it a crime to belong to a reference class? *Journal of Political Philosophy*, 9(2), 168–181.
- Cresswell, M. J. (1976). The semantics of degree. In *Montague grammar* (pp. 261–292). Elsevier.
- Degen, J., Tessler, M. H., & Goodman, N. D. (2015). Wonky worlds: Listeners revise world knowledge when utterances are odd. In *Proceedings of the 37th annual conference of the cognitive science society*.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ebeling, K. S., & Gelman, S. A. (1988). Coordination of Size Standards by Young Childre. *Child Development*, 59, 888–896.
- Ebeling, K. S., & Gelman, S. A. (1994, aug). Children’s Use of Context in Interpreting "Big" and "Little". *Child Development*, 65(4), 1178–1192. Retrieved from <http://doi.wiley.com/10.1111/j.1467-8624.1994.tb00811.x> doi: 10.1111/j.1467-8624.1994.tb00811.x
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Franke, M., Dablander, F., Scholler, A., Bennett, E., Degen, J., Tessler, M. H., ... Goodman, N. D. (2016). What does the crowd believe ? A hierarchical approach to estimating subjective beliefs from empirical data. In *38th annual meeting of the cognitive science society*.
- Futrell, R., Wilcox, E. G., Morita, T., & Levy, R. (2018). Rnns as psycholinguistic subjects: Syntactic state and grammatical dependency. *arXiv preprint arXiv:1809.01329*.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*.
- Goodman, N. D., & Stuhlmüller, A. (2014). *The Design and Implementation of Probabilistic Programming Languages*. <http://dippl.org>. (Accessed: 2019-9-19)
- Hájek, A. (2007). The reference class problem is your problem too. *Synthese*, 156(3), 563–585.

- Kemp, C., Shafto, P., & Tenenbaum, J. B. (2012). An integrated account of generalization across objects and features. *Cognitive Psychology*, 64(1-2), 35–73.
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1–35. doi: 10.1007/s10988-006-9008-0
- Kennedy, C., & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 345–381.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and philosophy*, 4(1), 1–45.
- Lassiter, D., & Goodman, N. D. (2013). Context, scale structure, and statistics in the interpretation of positiveform adjectives. *Proceedings of SALT*, 23(1), 587610.
- Lassiter, D., & Goodman, N. D. (2017). Adjectival vagueness in a Bayesian model of interpretation. *Synthese*.
- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge university press.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press.
- Linzen, T., Dupoux, E., & Goldberg, Y. (2016). Assessing the ability of lstms to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4, 521–535.
- Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122(3), 280–291.
- Qing, C., & Franke, M. (2014a). Gradable adjectives, vagueness, and optimal language use: A speaker-oriented model. *Proceedings of SALT*, 24.
- Qing, C., & Franke, M. (2014b). Meaning and Use of Gradable Adjectives: Formal Modeling Meets Empirical Data. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*.
- Qing, C., Goodman, N. D., & Lassiter, D. (2016). A rational speech-act model of projective content. In *Proceedings of the 38th annual conference of the cognitive science society*.
- Reichenbach, H. (1949). *The theory of probability. an inquiry into the logical and mathematical foundations of the calculus of probability*,. University of California Press.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive psychology*, 7(4), 573–605.
- Schmidt, L. A., Goodman, N. D., Barner, D., & Tenenbaum, J. B. (2009). How Tall Is Tall? Compositionality, Statistics, and Gradable Adjectives. In *Proceedings of the 31st annual conference of the cognitive science society*.
- Schöller, A., & Franke, M. (2017). Semantic values as latent parameters: Testing a fixed threshold hypothesis for cardinal readings of few and many. *Linguistic Vanguard*.
- Scontras, G., Tessler, M. H., & Franke, M. (2018). *Probabilistic language understanding: An introduction to the rational speech act framework*. Retrieved 2018-1-9 from <https://michael-franke.github.io/probLang>.
- Solt, S. (2009). Notes on the Comparison Class. In *International workshop on vagueness in communication*.
- Solt, S., & Gotzner, N. (2012). Experimenting with degree. *Proceedings of SALT*, 22, 166–187.

- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological review*, 124(4), 410.
- Tessler, M. H., & Franke, M. (2018). Not unreasonable: Carving vague dimensions with contraries and contradictions. In *Cogsci*.
- Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological review*, 126(3), 395.
- Tessler, M. H., Lopez-Brau, M., & Goodman, N. D. (2017). Warm (for winter): Comparison class understanding in vague language. In *Proceedings of the 39th annual conference of the cognitive science society*.
- Wilcox, E. G., Futrell, R., & Levy, R. (2021). Using computational models to test syntactic learnability. *lingbuzz preprint: lingbuzz/006327*.

Supplemental Information

Warm (for Winter): Inferring comparison classes in communication

Michael Henry Tessler^{1,2} and Noah D. Goodman^{2, 3}

¹Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

²Department of Psychology, Stanford University

³Department of Computer Science, Stanford University

Here we present additional information about methods, results, and cognitive models. Readers who are interested in the model code and analysis code are encouraged to consult the associated online repository: <https://github.com/mhtess/comparison-class-paper>.

Overview of Tasks

Our experimental procedures involve three coordinated experiments to (1) elicit test stimuli, (2) measure comparison class inferences and (3) measure adjective endorsements to facilitate model comparison (Figure S1). In Task 1, we empirically elicit test stimuli by having participants fill out phrasal templates that elicit sets of categories at the same level of abstraction, which differ in the general expectations participants have about those categories (e.g., basketball players [generally tall], jockeys [generally short], soccer players [sometimes tall and sometimes short]). From this set, we curate a set of experimental stimuli which we use in the other two tasks. Task 2 is the comparison class inference experiment described in the main text; this task is designed to test the qualitative predictions shown in Figure 1E (main text). Task 3 (an adjective endorsement, or truth judgment task) provides additional linguistic judgments that constrain model parameters to test the quantitative predictions of the comparison class inference model and quantitatively arbitrate between alternative models.

Experimental Methods

Bot check

In all tasks, participants were required to pass a simple language comprehension test that we designed in order to weed out bots and other bad-faith participants. The test involved a sentence in which a named speaker (e.g., Joseph) says to a named listener (e.g., Elizabeth) “It’s a beautiful day, isn’t it?”. Participants were asked to type in a text box to whom the speaker (in this case: Joseph) is talking (i.e., Elizabeth). Speaker and listener names were randomized in a way that could not be read off the source .html file. Participants

were given three attempts to correctly identify the listener. If they did not succeed within 3 attempts, they would be unable to proceed with the experiment. Since participants who fail this test are not allowed to proceed with the experiment, we do not have a count of how many participants fail this check.

Stimuli Generation Task (Task 1)

In this task, participants ($n = 50$) filled out phrasal templates for adjective pairs (e.g., *big* and *small*), in which three missing noun phrases which appeared in the grammatical subject of the sentence were described as either generally having one adjective apply to them (e.g., _____ *are generally big*; _____ *are generally small*) or as sometimes having either adjective apply to them (e.g., _____ *are sometimes big and sometimes small*). Participants filled out one template for each of 15 pairs of adjectives that describe physical dimensions (Table 1).

From this set of 750 responses, we curated a collection of 90 “item sets”, each of which consists of 3 relatively subordinate level categories that differ in their general expectations about the degree (e.g., Winter, Spring, and Summer, which differ in their general expectations about the typical temperature). In addition, each of these item sets is associated with a common, relatively superordinate level categories (e.g., days of the year).

Comparison Class Inference (Task 2)

In this experiment, we measure comparison class inferences by having participants rephrase a speaker’s statement which involves a scalar adjective in a way that makes the comparison class explicit. We measure comparison class inferences using a free-production measure to provide further ecological validity to our measurements of comparison classes: We wish to see if listeners spontaneously adjust their comparison class depending on world knowledge and pragmatic reasoning. A smaller-scale, forced-choice version of this task was reported in Anonymous (2017). Sample size, exclusion criteria, regression analysis, and cognitive model analysis were preregistered: osf.io/xuc96.

Participant restrictions and exclusion criteria. We recruited 837 participants from Amazon’s Mechanical Turk. Participants were restricted to those with U.S. IP addresses with at least a 95% work approval rating. In addition, participants were required to pass a simple language comprehension test that we designed in order to weed out bots and other bad-faith participants, described above.

In addition to the above restrictions, we excluded participants based on both a task comprehension question appearing before the main trials and a memory check trial appearing after the main trials. In the task comprehension / warm-up trial, participants were told they they would be asked to rephrase something a person said: the person said a word that is relative and their task was to figure out what the word was relative to. They were given the example of *John says: “The Empire State Building is tall”* and asked to fill-in a sentence with the same kind of response they would do on the main trials (i.e., *The Empire State Building is tall relative to other _____*). Participants were told to fill in the blank with a group or category that makes the most sense and to use their common sense. Responses to this warm-up trial were used as a basis for exclusion (any response other than buildings, structures, towers, skyscrapers, or any misspellings thereof). Invalid responses

Figure S1. Overview of Experimental Tasks. Task 1: Using a structured production task, we elicit sets of stimuli that all share the feature of containing categories generally judged as having either a positive or negative adjective (X , Y ; e.g., big or small) applied to them as well as a control category (Z ; e.g., sometimes big and sometimes small). The task is designed in a way to elicit three categories of the same basic- or superordinate-level category (W). Task 2: Free-production task to elicit the comparison class. Task 3: Forced-choice task where participants judge whether a member of the subordinate level category would be judged as having the adjective applied explicitly relative to the basic/superordinate level category. This task serves to provide additional data to constrain the parameters of the comparison class inference model.

to the warm-up trial were most often indicative of copying some part of the text on the screen and pasting it into the response box (e.g., responding with the name of the speaker, just putting the adjective “tall”, or responding with a whole sentence without a comparison class “that is tall”). 55 participants were excluded for providing an invalid response to this task comprehension question.

After the main comparison class inference trials, participants completed a memory check trial asking which adjective–Noun Phrase (NP) combinations appeared on the main trials. At the end of the task, participants were asked a memory check question where they had to select, from a list of 10 options, all of the items they could recall seeing. In the memory check, items were shown as adjective – noun pairs (“tall – basketball player”) and the 5 distractors were either color or multidimensional adjectives paired with a category that was not used in our test stimuli (e.g., “green – tennis ball”; “beautiful – painting”). Participants were excluded if they answered fewer than 7 out of 10 memory check questions correctly. A total of 59 participants failed this check, though 27 of these also failed the task comprehension check.

A total of 87 participants were excluded for meeting at least one of these exclusion

criteria, leaving a sample of 750 participants for the primary analyses.

Adaptive data collection procedure. Our experiment contains 540 unique items (adjective – NP pairs), which are highly heterogeneous. In addition to testing the main qualitative predictions of our models (adjective polarity X category expectations interactions), this experiment was designed to elicit high inter-item variability in comparison class inferences. As a result, we expect some items to exhibit low intra-item variability (i.e., all or almost all participants respond with the same comparison class); for these items, we would require relatively fewer data points to estimate the parameter of interest – the probability of a subordinate vs. superordinate comparison class – since it would either be close to ceiling or close to floor.

We thus used a sequential sampling method that we deployed on an item-wise basis, wherein we paused data collection after collecting 35 responses for each of the 540 items. We then analyzed the partial data set on a by-item basis to see which (if any) items had received exceedingly consistent responses, which we defined to be at least 33 out of 35 ($> 94\%$ agreement) of the same response. For those items that received exceedingly consistent responses, we stopped data collection at these 35 responses. We then continued collecting data on the items that received variable responses, until we had data from 750 participants after exclusion criteria were applied. This procedure allowed us to focus resources on providing better estimates for the items with more intra-item variability in responses. This adaptive procedure was decided ahead of time and is documented in the pre-registration report.

Response preprocessing. Our text preprocessing procedure was divided into two stages: Exclusions and Corrections.

Exclusions. We excluded responses that did not make sense given the context. These included but were not limited to responses that were simply the adjectives (e.g., *tall*), the names of the characters (e.g., *Alex*), and copied portions of the text on the screen (e.g., the sentence prompt: *the street is wide*). These invalid responses totaled 443 in total (1.6%).

We additionally removed from the data set responses that seemed to be the result of a failure of the participant to grasp the intended referent of the utterance. For example, in the context sentence: “Alexander is in the forest and hears a woodpecker. Alexander says, *It is loud*”, the intended referent of the pronoun *it* is the woodpecker; however, some participants responded *it is loud relative to other forests*, which suggests they understood the referent to which *it* refers to be the forest. This kind of “reference failure” responses numbered 267 in total (or, 1.0% of the remaining responses).

Corrections. In order to improve alignment across responses, we corrected for misspellings (462.0; 1.8%), lemmatized to align plural and singular responses (e.g., *baby* and *babies*; 3823.0; 14.5%), and replaced some responses with obvious synonyms given the context (e.g., *child-kid*; *booze-alcoholic drink*; *humans/individuals-people*; 660.0; 2.5%).

Response analysis. After preprocessing the raw text responses, we performed an automatic analysis of the comparison class responses by checking whether or not the preprocessed responses contained the subordinate NP presented in the experiment (e.g., *basketball player*) as a substring.

Pilot testing suggested that participants primarily provide comparison class phrases that are identical to the subordinate noun phrase by which the referent is introduced (*subordinate-NP*, e.g., a basketball player) or a more superordinate category (*superordinate-*

Figure S2. Proportion of responses under a three-way classification: mention subordinate, mention the modal superordinate response, or mention another, relatively superordinate category. For negative polarity adjectives, the proportion of both the “superordinate” and “other” responses increases as general expectations lower (and thus become more consistent with the adjective polarity; e.g., a short gymnast). For positive polarity adjectives, we see the reverse trend. Error-bars denote bootstrapped 95% CIs.

NP). In order to convey a category at least as specific as the subordinate-NP, one must include the subordinate-NP in the response (e.g., *male basketball players* is more specific than *basketball players* and includes the substring *basketball player*). One exception to this generality is utterances involving negation markers (e.g., “all kinds of people except basketball players”), but no responses in our data set involved negation in this way. Thus, we automatically categorized preprocessed responses as either subordinate or superordinate by checking whether the preprocessed response contained the subordinate-NP as a substring.

Supplementary data analysis.

Distribution of response types across experimental factors. In the main text, we reported that roughly 60% of responses explicitly mentioned the subordinate category, 25% of responses explicitly mentioned the modal superordinate category, and 15% of responses referred to some other, relatively superordinate category. We examine this 3-way classification of responses by the experimental factors that are the focus of this paper: general expectations about the category and the adjective polarity. Figure S2 shows the distribution of response types by general expectations and adjective polarity. Overall, we find a pattern of responses that is consistent with the idea that when the adjective is consistent with the general expectations about the category, a listener will look for a more superordinate comparison class.

Regression on subset of responses that explicitly mention subordinate or superordinate-NP. In the main text, we analyze the full data set of preprocessed responses. Here, we report statistics on a subset of the data set in which the participants preprocessed responses explicitly mention either the subordinate-NP (the NP by which the referent was introduced) or the superordinate-NP (an NP that the authors pre-specified as

the superordinate-NP for a set of three subordinate-NPs). 4007 responses (15.2%) are thus excluded from this analysis.

The results on the subset of responses that explicitly mention either the subordinate or superordinate-NP are entirely consistent with the results derived from the full data set. When the subordinate category was expected to be near the high-end of the scale (e.g., basketball player), the positive-form adjective (e.g., tall) led to fewer subordinate comparison classes than the negative-form adjective (e.g., short) in comparison to the control, middle-of-the-scale items (e.g., soccer player): posterior mean beta-weight and 95% Bayesian credible interval: $\beta = -1.51 [-1.94, -1.10]$. Again, this interaction was the result of the high-end-of-the-scale subordinate categories showing more subordinate comparison class inferences for the negative adjective (e.g., short) than for the positive adjective (e.g., tall): $\beta = 1.49 [1.12, 1.86]$, a preference which was not observed for the middle-of-the-scale, control items (e.g., *soccer players*; $\beta = 0.02 [-0.23, 0.28]$). As before, a comparable interaction was observed for categories that were expected to be near the low-end of the scale (e.g., jockey): Hearing the positive-form adjective led to credibly more subordinate inferences than hearing a negative-form adjective, in comparison to the middle-of-the-scale subordinate categories: $\beta = 1.42 [1.02, 1.83]$. Again, this was driven by the behavior of the low-end-of-the-scale subordinate categories, which showed a stronger preference for the superordinate comparison classes for negative than for positive adjectives ($\beta = -1.44 [-1.87, -1.04]$).

The inferences measured from this reduced data set are similarly highly symmetric across the low-end vs. high-end of the scale subordinate categories. We observe an overall preference for subordinate comparison classes for the control (middle-of-the-scale) subordinate categories (e.g., *soccer players*; $\beta = 1.68 [1.16, 2.21]$) and no overall differences in this preference for items at the high-end of the scale (e.g., *basketball players*; $\beta = 0.10 [-0.25, 0.44]$) or the low-end of the scale (e.g., *gymnasts*; $\beta = 0.23 [-0.15, 0.61]$). The inferences that result from adjectives that are in conflict with a listener’s general expectations of the categories (e.g., tall gymnasts vs. short basketball players) were not different between the low-end and high-end items ($\beta = 0.11 [-0.42, 0.66]$), nor were the inferences from adjectives that were consistent with general expectations about a category (e.g., short gymnasts vs. tall basketball players; $\beta = 0.15 [-0.35, 0.65]$).

Universally consistent responses. 16 of our 540 items (adjective – noun pairs) had universally consistent lemmatized responses wherein all participants gave the same response; this number increases to 89 on the restricted data set that only includes responses that match either the subordinate or superordinate responses (described in the previous section). Of these universal consistent responses (i.e., responses that were at floor or at ceiling), all but one were cases where 100% of participants gave a subordinate-NP response (i.e., a response that was the same as the NP with which the referent was introduced). For example, 100% of valid responses for a “loud church” all concerned the subordinate category *church*. The lone NP-adj pair that led to all superordinate responses (for either data set) was the “cheap plastic bracelet”, which all participants said was cheap relative to other *bracelets*. Intriguingly, the numeric counts of items that received 100% subordinate-NP responses mirrors the overall pattern of responses collapsed across items shown in the main text: adjectives that conflict with the general expectations about the category (e.g., a short basketball player) are more likely to receive subordinate-NP comparison class responses

Figure S3. Comparison Class Inference experimental results removing responses that mention neither subordinate nor superordinate categories. Proportion of paraphrases that contained the Subordinate NPs (e.g., *basketball player*) with which the referent was introduced, as a function of the general expectations listeners have about the category (x-axis) and the polarity of the adjective used to describe the category (e.g., *tall* = positive, *short* = negative). Bars represent overall means and error bars are bootstrapped 95% confidence intervals. Each dot represents the mean of a single item and lines connect Subordinate NPs described with different adjectives (e.g., *tall* and *short* basketball player). Dots are jittered horizontally to improve visual clarity.

(Table S1 for counts of universally consistent responses from the full data set; Table S2 for universally consistent responses from the restricted data set).

Adjective Endorsement Task (Task 3)

This experiment served to generate further data that could be used to constrain the model parameters shown in Figure S7. Sample size, sampling procedure, exclusion criteria, and data analysis were preregistered: osf.io/vdkbp.

Participants. We recruited 400 participants from Amazon’s Mechanical Turk, restricted to those with U.S. IP addresses with at least a 95% work approval rating. This number was arrived at with the goal of estimating endorsement probabilities for each item with a 95% confidence interval width of 0.2 for each item. Because our measure is a two-alternative forced-choice and because we expect some items to receive fairly categorical judgments (i.e., low between-participant variability), we employed an analogous sequential sampling procedure to that employed in Task 2 (Comparison Class Inference, see above) in order to over-sample items that exhibit high between-participant variability. We first collected 35 responses for each item, computed 95% CIs for each item, and stopped collecting data for items whose 95% CIs were smaller than 0.2. After the first 35 responses, we collected more responses for each item until the 95% CI for that item was smaller than 0.2 or we reached our predetermined sample size. The experiment took about 10 minutes and

Adjective Polarity	low	medium	high
negative	1	2	4
positive	7	0	1

Table S1

Counts of items that had 100% subordinate-NP comparison class responses, as a function of the adjective polarity (positive vs. negative) and general expectations about the category (low, medium, high) computed over the full data set.

Adjective Polarity	low	medium	high
negative	12	11	21
positive	23	15	7

Table S2

Counts of items that had 100% subordinate-NP comparison class responses, as a function of the adjective polarity (positive vs. negative) and general expectations about the category (low, medium, high) computed over the restricted data set.

participants were compensated \$1.50 for their work.

Procedure. On each trial, participants were given a sentence introducing a member of a subordinate category (e.g., *You step outside during the winter.*). This was followed by a question asking if the participant would endorse the positive- and/or negative-form adjective explicitly relative to the superordinate category (e.g., *Do you think it would be warm relative to other days of the year? Do you think it would be cold relative to other days of the year?*). Both questions appeared on the screen on the same trial. Participants could respond to each question with either a *yes* or *no* judgment (2 judgments per trial). Each participant completed 48 items.

As with the comparison class inference experiment, after the main trials, participants completed a memory check trial asking which adjective–NP combinations appeared on the main trials. Participants were asked a memory check question where they had to select, from a list of 10 options, all of the items they could recall seeing. In the memory check, items were shown as adjective – noun pairs (“tall – basketball player”) and the 5 distractors were either color or multidimensional adjectives paired with a category that was not used in our test stimuli (e.g., “green – tennis ball”; “beautiful – painting”). Participants were excluded if they answered fewer than 7 out of 10 memory check questions correctly. A total of 103 participants failed this check, leaving 297 participants for the main analysis.

Materials. The experimental materials were the same as used in Task 2 (Comparison Class Inference). Each trial contained 2 judgments (positive and negative adjectives).

Results. The primary goal of this experiment is to validate the stimuli generated in Task 1 and to use the data to constrain the parameters governing world knowledge in the joint Bayesian data-analytic model. We build a Bayesian logistic mixed-effects model to predict participants’ responses as a function of the general expectations about the subordinate category (low, medium, high; dummy coded with the medium category as the reference level), the adjective (positive vs. negative; difference coded), and their interaction; in addition, we include the maximal mixed effects structure by-item set and by-participant that mirrors this fixed effects structure.¹

The pattern of endorsements was highly consistent with the a priori general expectations by which the stimuli were generated (Figure S4). When the subordinate category was

¹The model is $\text{judgment} \sim \text{gen\_expectations} * \text{adjective\_polarity} + (1 + \text{gen\_expectations} * \text{adjective\_polarity} \mid \text{participant}) + (1 + \text{gen\_expectations} * \text{adjective\_polarity} \mid \text{item\_set})$.

Figure S4. Adjective Endorsement (Task 3) experimental results. Proportion of sentence endorsements as a function of the general expectations listeners have about the category (x-axis) and the polarity of the adjective used to describe the category (colors; e.g., *tall* = positive, *short* = negative). Bars represent overall means and error bar is a bootstrapped 95% confidence interval. Each dot represents the mean of a single item and lines connect Subordinate NPs described with different adjectives (e.g., *tall* and *short* basketball player). Dots are jittered horizontally to improve visual clarity. The biggest outlier for the low general expectations was the loudness of the *owl*, which was judged to be more loud than quiet; the biggest outlier for the high general expectations was the fabric of *wool*, which was judged to be more light than heavy.

expected to be near the high-end of the scale (e.g., basketball player), the positive-form adjective (e.g., tall) was endorsed much more strongly than the negative adjective (e.g., short) in comparison to the control, middle-of-the-scale items (e.g., soccer player): posterior mean beta-weight and 95% Bayesian credible interval: $\beta = 5.72$ [4.81, 6.64]. This interaction was the result of the high-end-of-the-scale subordinate categories showing more endorsements for the positive (e.g., tall) than for the negative (e.g., short) adjective: $\beta = -5.38$ [-6.02, -4.75]. The same endorsement profile was not observed for the middle-of-the-scale, control items (e.g., soccer players, for which the positive and negative adjective were endorsed equally as often $\beta = -0.34$ [-1.04, 0.35]). A comparable interaction was observed for categories that were expected to be near the low-end of the scale (e.g., jockey): The positive-form adjective was endorsed credibly less than the negative-form adjective, in comparison to the middle-of-the-scale subordinate categories: $\beta = -4.58$ [-5.32, -3.84]. Again, this was driven by the behavior of the low-end-of-the-scale subordinate categories, which showed strong endorsements for the negative than for the positive adjectives ($\beta = 4.92$ [4.22, 5.62]).

The results from the endorsement task were remarkably symmetric across the low-end vs. high-end of the scale subordinate categories. We observe for the control (middle-of-the-scale) subordinate categories (e.g., soccer players), a slight preference to reject both adjectives over accepting both (i.e., *neither tall nor short* was a more common pattern than

Figure S5. Task 3 (Adjective Endorsement) results for a subset of the items.

both tall and short; $\beta = -0.12 [-0.21, -0.02]$). The mean endorsements, averaging over positive and negative adjectives, for items at the high-end of the scale (e.g., *basketball players*) was not different from chance ($\beta = 0.01 [-0.09, 0.11]$) nor were the average endorsements of adjectives for categories at the low-end of the scale (e.g., *gymnasts*; $\beta = -0.02 [-0.11, 0.07]$). Furthermore, the endorsements of adjectives that are in conflict with a listener's general expectations of the categories (e.g., tall gymnasts vs. short basketball players) were not different between the low-end and high-end items ($\beta = 0.20 [-0.25, 0.65]$), nor were the inferences from adjectives that were consistent with general expectations about a category (e.g., short gymnasts vs. tall basketball players; $\beta = -0.26 [-0.71, 0.18]$).

Alternative Models

The inference about the comparison class outlined in the main text involves a listener reasoning about a speaker reasoning about a listener. The full Rational Speech Act model is given by these equations (reproduced from the main text):

$$L_1(x, c \mid u; k) \propto S_1(u \mid x, c) \cdot P(x \mid k) \cdot P(c \mid k) \quad (S1)$$

$$S_1(u \mid x, c) \propto \exp(\alpha \cdot (\ln L_0(x \mid u, c) - \text{cost}(u))) \quad (S2)$$

$$L_0(x \mid u, c) = \int_{\theta} L_0(x, \theta \mid u, c) d\theta \quad (S3)$$

$$L_0(x, \theta \mid u, c) \propto \delta_{[u](x, \theta)} \cdot P(x \mid c) \cdot P(\theta) \quad (S4)$$

Figure S6. Alternative model predictions. A: A literal model that does not represent a speaker’s representation of the context as separate from their own. This listener effectively answers the question of what is more likely: a basketball player who is tall or a person who is tall? This model predicts the exact opposite pattern of data from the pragmatic model reported in the main text. B: A more sophisticated pragmatic listener model than that reported in the main text. This model reasons about which comparison class a speaker believes that the listener would infer. This model produces the same pattern of results as the simpler pragmatic model reported in the main text.

Literal alternative model

One might question whether the inference about the comparison class is necessarily a pragmatic inference which requires recursive reasoning. We can investigate this question by reformulating the comparison class inference spelled out in Equation S1 in terms of a literal listener model (*a la* Equation S4). This literal comparison class inference model is given by:

$$L_0(x, \theta, c \mid u, k) \propto \delta_{\llbracket u \rrbracket}(x, \theta) \cdot P(x \mid c) \cdot P(\theta) \cdot P(c \mid k) \quad (\text{S5})$$

Similar to the pragmatic listener model (Eq. S1), this listener can use their knowledge of the referent k to constrain the hypothesis space of comparison classes (e.g., with the knowledge that the referent is a basketball player, consider comparison classes that the same as or superordinate to the class of basketball players). Unlike the pragmatic listener model (Eq. S1), however, the literal listener version of the model does not hold different representations of the referent in mind: The pragmatic listener has their private representation of the referent—given by the prior distribution of the degree $P(x \mid k)$ —and imagines a speaker who acts assuming some comparison class— $S_1(u \mid c)$ —where c and k may or may not index the same class (e.g., the listener may know the referent is a basketball player – $k = \text{basketball players}$ – but believe the speaker was assuming a *person* comparison class – $c = \text{people}$). The literal listener version of the model has no way of separating these representations. In effect, this listener is answering a slightly different question from the comparison class inference problem. The question this alternative model is answering is: what is more likely—a basketball player whose height is greater than some threshold or a person whose height is greater than some threshold? This alternative model predicts the exact opposite pattern of results from the pragmatic listener model (Figure S6A).

Alternative pragmatic model

The pragmatic comparison class inference listener model (Eq. S1) reasons about which comparison class a speaker is more likely to be assuming. That is, the speaker (Eq. S2) is presumed to be assuming that a particular comparison class is already in the common ground, analogous to a presupposition (e.g., saying “My car is in the shop” presupposes that the speaker owns a car). Speakers may be aware, however, that the comparison class is not in the common ground; still, they may still avoid articulating a comparison class (e.g., uttering “He’s short for a basketball player”) if the listener can reasonably be assumed to infer the comparison class.

This kind of inference is more sophisticated: It involves a listener reasoning about the comparison class that a speaker believed the listener would infer. As demonstrated above with the Literal Alternative model, this first-order listener inference about the comparison class itself involves pragmatic reasoning. This higher-order pragmatic reasoning model is given by the following equations:

$$L_2(x, c \mid u, k) \propto S_2(u \mid x, c) \cdot P(x \mid k) \cdot P(c \mid k) \quad (\text{S6})$$

$$S_2(u \mid x, c) \propto \exp(\alpha \cdot (\ln L_1(x, c \mid u) - \text{cost}(u))) \quad (\text{S7})$$

where L_1 is the pragmatic inference model defined above and in the main text (Eq. S1). The primary difference between this model and the model given by Eqs. S1 and S2 is in the speaker S_2 (Eq. S7) (vs. the S_1 shown in Eq. S2). S_2 chooses their utterance by taking into account the fact that the listener L_1 (Eq. S1) is uncertain about the comparison class. As shown in Figure S6B, this more sophisticated pragmatic inference model arrives at the same conclusions about the likely comparison class given different general expectations about the category and the adjective heard. Further, the inferences of this model, like those of the simpler pragmatics model, are resilient to reasonable choices of alternative utterances; most notably, if the set of alternative utterances provides a way to explicitly articulate the comparison class (e.g., the speaker could have said *They’re tall for a basketball player*), the same inferences result from hearing the utterance without a comparison class.

Since this alternative pragmatic model derives the same qualitative inference for the contexts we study about the comparison class as the model presented in the main text, we do not compare the two models in the Quantitative Analysis section. It is likely that the predictions of these two models would come apart under different contextual manipulations of speaker and/or listener knowledge. For example, the model presented in the main text might be better suited for third-party conversations and its architecture could be useful for developing models of word learning from overheard speech (Akhtar, Jipson, & Callanan, 2001; Martínez-Sussmann, Akhtar, Diesendruck, & Markson, 2011). On the other hand, the more sophisticated pragmatic inference model presented in this supplement might be better suited for dyadic interactions with informed and intentional interlocutors.

Bayesian Data Analysis

We generate quantitative predictions for the Comparison Class Inference by constructing a joint Bayesian data-analytic model that synthesizes the data from the Comparison

Figure S7. Joint Bayesian data analytic strategy of the maximal model. Two related RSA models (the comparison class inference model L_1 and the adjective endorsement model S_1) directly predict the data from Tasks 2 & 3 (d^{cci} and d^{adj}), respectively. Each of these models relies upon world knowledge, which varies by the degree d (e.g., height) and category k (e.g., basketball players)— $P(x^{d,k})$, assumed to be Normal distributions with unknown mean μ_k^d and variance σ_k^d . The prior probability of a comparison class $c_{d,k}$ is used only in the L_1 model, and is assumed to be a logistic linear function encoding a basic level bias β_0 and an effect of frequency of the noun phrase $\hat{f}_k$ given by β_1 . The referent category of each stimulus may be either a subordinate level category or a basic-level category, determined by parameter $z_{d,k}$, which gates between using two different basic-level bias comparison class parameters. Finally, each RSA model has its own speaker optimality parameter α .

Class Inference experiment (Task 2) and the Adjective Endorsement experiment (Task 3) using Rational Speech Act models to predict the data from each experiment. The RSA model for the comparison class inference data is described in the main text of the experiment and is given by Eq. S1.

We model the adjective endorsement data using the pragmatic speaker component S_1 (Eq. S2) of the comparison class inference model, with three slight modifications: (1) the comparison class c is assumed to be the superordinate comparison class, since the endorsement task concerns sentences drawing explicit comparisons to a superordinate class,

(2) the speaker has two alternative utterances: produce the adjective in question (e.g., tall or short) or stay silent (following Degen & Goodman, 2014; Tessler & Goodman, 2019), and (3) the speaker’s utility is computed via an expectation about the degree value x for the referent given by its category membership k : $x \sim P_k$ (e.g., the speaker does not have a particular height in mind to convey but averages over their expectations about the height of the referent). The adjective endorsement RSA model we use is given by:

$$S_1(u \mid k) \propto \exp(\alpha_2 \cdot \mathbb{E}_{x \sim P_k} \ln L_0(x \mid u)) \quad (\text{S8})$$

where $L_0(x \mid u)$ is Equation S4 printed above.

Each RSA model has a global speaker optimality parameter α_i which determines the degree to which speakers are assumed to be informative. This parameter is not of direct theoretical interest and we use priors consistent with the previous literature on RSA models: $\alpha_i \sim \text{Uniform}(0, 20)$.

Degree priors (World knowledge)

The two RSA models (Eq. S1 for the Comparison Class Inference task and Eq. S8 for the Adjective Endorsement task) depend upon the same world knowledge priors $P(x \mid k)$ (e.g., expectations about the height of basketball players), which vary as a function of the subordinate category k . Thus, we put uncertainty over these world knowledge prior distributions and infer the priors that best jointly explain the data from the two tasks.

Only the relative values for $P(x \mid c = c_{sub})$ and $P(x \mid c = c_{super})$ affect model predictions. We further assume the degree priors follow Gaussian distributions. Hence we fix each superordinate distribution to be a standard normal distribution $P(x \mid c = c_{super}) = \mathcal{N}(0, 1)$ and the subordinate priors are Gaussian distributions with unknown means and variances $P(x \mid c_{sub}) = \mathcal{N}(\mu_{sub}, \sigma_{sub})$; the subordinate priors thus have standardized units. We put priors over these mean and variance parameters of each subordinate category distribution k for each degree d : $\mu_k^d \sim \text{Uniform}(-3, 3)$, $\sigma_k^d \sim \text{Uniform}(0, 3)$, and infer their likely values using Bayesian Data Analysis (left-hand side of Figure S7).

Comparison class priors and description of alternative models

The comparison class prior $P(c)$ reflects listeners’ expectations of what comparison classes are likely to be used in conversation. Baseline expectations about conceptual comparison classes could be a function of the level of abstraction of the categories in question as well as the usage frequency of the noun phrases used to describe those categories. For example, basic-level categories may be more probable conceptual comparison classes because of their utility in everyday reasoning (Rosch & Mervis, 1975); additionally, we might expect the relative probability of basic-level vs. subordinate level categories to differ from basic vs. superordinate categories. To investigate these possibilities, we construct and compare models which differ in how the comparison class prior is parameterized.

Flat prior model (baseline). Our baseline model makes inferences based only on distributional knowledge about scalar properties of categories $P(x \mid k)$ (described above). For this model, the comparison class prior is uninformed and does not provide an *a priori* preference for the relatively subordinate or relatively superordinate comparison class; thus,

all of the by-item variability in comparison class inferences in this model must be explained as a function of the distributional knowledge about properties.

Basic-level bias. The *basic-level bias* model assumes some *a priori* preference for the Subordinate or Superordinate comparison classes. We parameterize the comparison class prior via a logistic model, where a basic-level bias plays the role of an intercept term β_0 : $P(c) = \text{logistic}(\beta_0)$. Our crowd-sourced stimuli generation procedure (Task 1) presents interesting inferential challenges for this analysis. *A priori*, we do not know if the NPs we use to introduce the referents (Referent NPs; those generated by participants in Task 1) are basic-level categories or subordinate level categories, and hence, whether the more general comparison class would correspond to a superordinate-level category or a basic-level category. A basic-level bias could plausibly operate differently for a subordinate vs. basic-level inference than for a basic vs. superordinate level inference. Specifically, superordinate comparison classes might be the most implausible, because superordinate categories are more heterogenous in comparison to basic-level or subordinate-level categories. Thus, we endow our data-analytic model with two regression coefficient parameters corresponding to the intercept term (i.e., the basic-level bias term), and introduce a Bernoulli random variable z for each NP to indicate whether it is a subordinate-level term or basic-level term (Figure S7).

Frequency effect. The *frequency effect* model assumes that the NPs that denote comparison classes which are more common in usage could have a higher *a priori* probability of serving as the comparison class. The comparison class prior in this model is thus a (logistic) function of the relative usage frequency f_k of the subordinate NP vs. superordinate NP estimated from the Google WebGram corpus²: $P(c) = \text{logistic}(\beta_1 \cdot \log(\frac{\hat{f}_{sub}}{\hat{f}_{super}}))$. We assume a usage-based frequency effect to operate independent of the level-of-abstractness of the category; hence, we do not infer separate β_1 's for the imputed subordinate vs. basic-level categories.

Full model: Basic-level bias and frequency effect. To test whether both a basic-level bias and frequency effect jointly contribute to the comparison class prior, our full model assumes both contribute via a logistic-linear model: $P(c) = \text{logistic}(\beta_0 + \beta_1 \cdot \log(\frac{\hat{f}_{sub}}{\hat{f}_{super}}))$.

Priors for comparison class parameters. The priors over the regression parameters (which also apply to the alternative models described above, if they include that parameter) are $\beta_0^0, \beta_0^1 \sim \text{Gaussian}(0, 2)$, $\beta_1 \sim \text{Uniform}(-3, 3)$. The binary level-of-abstractness parameter which governs the direction and magnitude of the basic-level bias comes from a uninformed prior: $z_k \sim \text{Bernoulli}(0.5)$.

Model implementation

We ran four different BDA models, corresponding to the different ways of parameterizing the comparison class prior: (1) Flat Prior model (assumes the comparison class prior is always 50/50 between subordinate and superordinate class); (2) Basic-level bias, or intercept only model (assumes a basic-level bias), (3) Frequency effect, or slope only model (assumes an effect of corpus frequency, but no basic-level bias), (4) Basic-level bias

²Corpus accessed via https://corpora.linguistik.uni-erlangen.de/cgi-bin/demos/Web1T5/Web1T5_freq.perl

(a) Human data and model predictions for items with small residuals (left 4) and largest residuals (right 4) for comparison class inference data. Top two rows show data and model predictions for the main Comparison Class Inference experiment; bottom two rows correspond to the Adjective Endorsements (norming study). The log ratio frequency of subordinate NP to superordinate NP is shown in brackets next to the x-axis labels; a more negative number corresponds to a stronger prior belief in the superordinate category as the comparison class.

(b) Imputed prior distributions over degrees for 8 of the 90 item sets. These distributions were generated from the Maximum A-Posteriori parameter values inferred by conditioning on the Comparison Class Inference and Adjective Endorsement data sets.

Figure S8. Quantitative modeling results for eight sets of items which show the lowest and highest residuals for the comparison class inference data.

and Frequency effect (slope and intercept). Figure S7 shows the plate diagram corresponding the BDA model for the full, slope and intercept model. We implemented the RSA and Bayesian data analysis models in the probabilistic programming language WebPPL (Goodman & Stuhlmüller, 2014) and performed inference by running 7 MCMC chains with 500,000 iterations each, discarding the first 250,000 for burn-in. Convergence was checked through visual inspection of the different chains to ensure similar conclusions would be drawn from each chain independently.

We estimated the marginal likelihood of the data under each model by using an Annealed Importance Sampling algorithm implemented in WebPPL (Neal, 2001). We collected 4 samples from the AIS algorithm, which was run 250,000 steps per sample.

Supplementary model results

Figure S8 shows model predictions and inferred parameters for four sets of items which for which the full model does well at predicting the comparison class inference data and four sets of items where the full model does not do well at predicting the comparison class inferences. Predictions in the Comparison Class Inference task (Fig. 8(a), top rows) can be triangulated by examining the predictions in the Adjective Endorsement task (Fig. 8(a), bottom rows) and the inferred world knowledge parameters (Fig. 8(b)).

References

- Akhtar, N., Jipson, J., & Callanan, M. A. (2001). Learning words through overhearing. *Child development*, 72(2), 416–430.
- Degen, J., & Goodman, N. (2014). Lost your marbles? the puzzle of dependent measures in experimental pragmatics. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 36).
- Goodman, N. D., & Stuhlmüller, A. (2014). *The Design and Implementation of Probabilistic Programming Languages*. <http://dippl.org>. (Accessed: 2019-9-19)
- Martínez-Sussmann, C., Akhtar, N., Diesendruck, G., & Markson, L. (2011). Orienting to third-party conversations. *Journal of Child Language*, 38(2), 273.
- Neal, R. M. (2001). Annealed importance sampling. *Statistics and computing*, 11(2), 125–139.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive psychology*, 7(4), 573–605.
- Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological review*, 126(3), 395.