



Latent scope bias in categorization



Abigail B. Sussman^{a,*}, Sangeet S. Khemlani^b, Daniel M. Oppenheimer^c

^a University of Chicago Booth School of Business, 5807 S. Woodlawn Ave., Chicago, IL 60637, USA

^b Intelligent Systems Section, Code 5515, Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, 4555 Overlook Ave. SW, Washington, DC 20375, USA

^c UCLA, Anderson School of Management, 110 Westwood Plaza, Los Angeles, CA 90095, USA

HIGHLIGHTS

- We investigate how people categorize exemplars given incomplete information.
- We define scope as the number of distinct features category membership implies.
- Results show bias for grouping exemplars in categories with narrower latent scope.
- Preferences extend to verbal and visual categorization tasks.

ARTICLE INFO

Article history:

Received 22 January 2013

Revised 30 July 2013

Available online 7 December 2013

Keywords:

Narrow scope bias

Categorization

Stereotyping

Person perception

Causal reasoning

ABSTRACT

Categories often have unobservable diagnostic features. For example, if a person is a lawyer, one might expect him to be both well dressed and knowledgeable about the law. However, without observing the person in a courtroom, one cannot tell whether or not he is knowledgeable about the law. How might we categorize the well-dressed person before we know whether or not he possesses a particular category feature? Two studies showed that, all else equal, individuals prefer to group exemplars into categories that specify fewer unobserved and unobservable features — i.e., those that have a narrower latent scope — to those with a broader latent scope. In Experiment 1, participants were more likely to classify novel exemplars as part of a social category that had a narrower latent scope in a verbal task. Experiment 2 demonstrated that the scope bias generalizes to contexts in which category structure is never explicitly specified.

© 2013 Elsevier Inc. All rights reserved.

Introduction

The process of generalizing knowledge from a known category to a novel instance is central to the way we perceive the world, and it has permeated intellectual debate since Plato (Statesman, 261e et seq.). The mechanisms by which we categorize individuals constrain key components of social perception, such as stereotyping, impression formation, and even recall of information about others (e.g., Cantor & Mischel, 1979; Cohen, 1981; Higgins, Rholes, & Jones, 1977; Klein, Loftus, Trafton, & Fuhrman, 1992; Macrae & Bodenhausen, 2000; Stangor, Lynch, Duan, & Glass, 1992; Tajfel, Billig, Bundy, & Claude, 1971), as well as broader aspects of judgment and decision-making (for a review, see Murphy, 2002).

Given that people fit into many different categories, one tradition in person perception research has attempted to discern which categories are used for prediction (e.g., Crisp & Hewstone, 2007; Kunda, Miller, &

Claire, 1990; Macrae, Bodenhausen, & Milne, 1995). Rather than making inferences from multiple possible categories, people tend to infer attributes based on the most likely category (Malt, Ross, & Murphy, 1995). It is therefore critical to understand how individuals determine the most likely category.

In laboratory studies on categorization, participants typically have complete information about which relevant features a putative category member possesses — participants are told that the member either possesses or does not possess a feature. However, this design is not paralleled in everyday life, where knowledge about an exemplar's features is frequently unknown or uncertain. The uncertainty complicates an already difficult categorization task. What strategies do people use to overcome informational limitations? Although categorization under uncertainty has received attention (e.g., Griffiths, Hayes, & Newell, 2012; Molden & Higgins, 2004; Murphy & Ross, 1994, 2005; Ross & Murphy, 1996; Verde, Murphy & Ross, 2005), there has been little study of how the structure of the category affects how people attempt to overcome missing or uncertain information.

Studies in explanatory reasoning suggest that people's knowledge of a category's causal structure may drive their categorization judgments. More than a decade of research has shown that both causal and explanatory reasoning play key roles in categorization (e.g., Ahn, Kim,

* Corresponding author at: University of Chicago Booth School of Business, 5807 S. Woodlawn, Chicago, IL 60637, USA.

E-mail addresses: asussman@chicagobooth.edu (A.B. Sussman), skhemlani@gmail.com (S.S. Khemlani), daniel.oppenheimer@anderson.ucla.edu (D.M. Oppenheimer).

Lassaline, & Dennis, 2000; Murphy & Medin, 1985; Lombrozo, 2009; Rehder, 2003a, 2003b; Rehder & Hastie, 2001; Sloman, Love, & Ahn, 1998; Waldmann, Holyoak, & Fratianne, 1995). If explanatory reasoning and categorization recruit the same cognitive mechanisms, then an examination of the processes that underlie the generation and evaluation of explanations might help account for performance on categorization tasks along with consequences for stereotyping and inference.

A recent analysis of the role of causal structure in explanatory reasoning explored how people determine the best explanation for a set of observations when information is incomplete (Khemlani, Sussman, & Oppenheimer, 2011). The researchers identified a narrow *latent scope* bias. Latent scope describes the number of effects for which an explanation could potentially account, regardless of whether or not the effects are observable. People appear to prefer explanations with narrower latent scope. For example, consider the following:

A causes X and Y.
 B causes X, Y, and Z.
 We observed X; no information is known about Y or Z.
 Which is more likely: A or B?

Khemlani et al. (2011) found that people prefer Explanation A in cases like this, because Explanation A causes fewer unobserved effects, even though these unobserved effects would not have been known even had they been present. In other words, when information is missing, people prefer explanations that make no predictions about items that are both unobserved and potentially unobservable (i.e., are not known to be either present or absent given the available evidence). In the context of categorization, the latent scope bias suggests that when the status of a set of features is unknown, people may prefer to align exemplars with categories that specify fewer features altogether, i.e., those that have narrower latent scope.

In this paper, we briefly review findings on the latent scope bias in explanatory reasoning and then provide empirical support for a parallel bias in the categorization of novel exemplars given limited information.

Latent scope in explanatory reasoning and categorization

The latent scope of an explanation can be thought of as the number of distinct effects for which the explanation can potentially account. An explanation's scope is *latent* because the possible effects that it can describe may not necessarily materialize or be observable. Explanations that could account for fewer effects have narrower latent scope than those that account for many effects. For instance, contrast two explanations for why someone might dye his hair and then shave his head: he dislikes his new hair color, or he is going through a mid-life crisis. The first explanation has narrower latent scope; it can only explain how someone might behave with respect to his hair. The second explanation has a broader latent scope: going through a mid-life crisis could also account for a wide range of other behaviors, many of which may never materialize. If we assume equal proportions of people who dislike their hair color and who are going through a mid-life crisis, then there is no normative reason to prefer one explanation to the other.

However, Khemlani et al. (2011) show that individuals exhibit a strong bias in favor of the explanation with narrower latent scope. The effect was robust even when base rates favored the broad scope explanation, and it persisted in more naturalistic domains in which the structure of the explanation was not made explicit. Khemlani et al. (2011) also ruled out several explanations for the effect, including the possibility that participants interpreted the absence of information about an effect to mean that the effect was not present.

In explanatory reasoning, latent scope is defined as the number of effects the explanation could account for. In categorization, we extend the notion of latent scope to refer to the number of distinct features that category membership implies. The scope is latent because not all features of an exemplar may be observable when inferring category

membership. For example, suppose that you observe someone who is knowledgeable about medications, and you have to decide whether she belongs to the category of pharmacists or physicians. If you do not know whether or not the person also has the ability to write prescriptions, and assume equal base-rates of pharmacists and physicians, would you be more likely to categorize the person as a pharmacist or a physician? The "pharmacist" category has narrower latent scope, since it specifies only one distinct features (e.g., is knowledgeable about medications), whereas the "physician" category has broader latent scope, because it specifies at least two distinct features (e.g., is knowledgeable about medications and has the ability to write prescriptions).

We next describe two studies that show that, parallel to the findings in explanatory reasoning, people prefer to place exemplars in categories with narrower latent scope. We first describe a test of the narrow latent scope hypothesis that used a verbal social categorization task before detailing an investigation of the bias in a visual category learning task. Both methods of examination reveal a robust narrow latent scope bias in categorization.

Experiment 1

Experiment 1 tested the narrow latent scope hypothesis by first presenting participants with verbal descriptions of a variety of social categories and category members. Participants then performed a classification task. Problems relied on fictional categories to ensure that variations in category structure, rather than prior knowledge, caused any observed variations in judgment. The narrow latent scope bias predicts that people should tend to believe that the exemplar is a member of the category that specifies a more limited set of features. The effect of the bias is accordingly to leave fewer unobserved features unknown.

Method

Participants

Forty-nine participants were recruited through Amazon.com's Mechanical Turk platform (for a discussion on the validity of results from this platform, see Paolacci, Chandler, & Ipeirotis, 2010) and participated in the study for monetary compensation.

Design and procedure

Participants were presented with a series of questions. In each question, participants received information about two otherwise unfamiliar categories. They were told explicitly that the categories had approximately equal numbers of category members. For example, participants were told about people who belonged to the Tokolo tribe:

In the jungles of the Amazon about half of the Tokolo tribe members are hunters, and the other half are spear fishermen. Both hunters and spear fishermen carry spears, but spear fishermen also carry nets.

In this example, the category of hunters is defined by fewer specific features (carrying spears) and is therefore considered to have narrower latent scope than the category of spear fishermen (who carry spears and nets). After participants were told about each category, they were informed about a specific person, for example, "You come across a tribesman who has a spear, but you don't know whether or not he also has a net." They were then asked to choose which category the person was more likely to belong to given the two alternatives.

Participants saw eight problems in total – four experimental and four control – each with a distinct set of categories (see Appendix A for problem forms and Appendix B for category descriptions). Across all problems, information about certain features was known and presented to the participant, and information about other features was unavailable. In the latter case, the fact that this information was unavailable was explicitly stated to the participant to clarify that the absence of information about a feature did not indicate the absence of the feature itself.

Experimental problems examined situations in which available information could not discriminate between the two categories (e.g., the tribesman has a *spear* in the case above). Control problems ensured that participants understood the task and would select broad scope categories when it was normative to do so (e.g., the tribesman has a *net*). The order of category names and whether each set of category descriptions was paired with a control or experimental question was counterbalanced; the order of questions and response options were randomized.

Results and discussion

Three participants were excluded prior to analysis because English was not their native language and one participant was excluded for taking the survey multiple times. Results, including significance levels, are consistent whether or not the participants' data were discarded. In response to control questions, participants judged unknown people more likely to be members of narrow latent scope categories only 10.6% of the time, indicating that they understood the task (one-sample subject-level Wilcoxon test relative to 50%, $z = 5.89, p < .001$). However, this pattern reversed for experimental problems: participants considered unknown people more likely to be members of the narrow latent scope categories 66.7% of the time when available information about the person was ambiguous (one-sample subject level Wilcoxon test relative to 50%, $z = 4.12, p < .001$).

Results show that the latent scope bias applies to social categorization. In this verbal task, participants were biased towards narrower latent scope categorizations when information was ambiguous. However, the categories and category members were defined through a limited set of features and explicitly stated rules governing category membership. Of course, in natural settings, people do not have access to explicit statements of the category structure. They have to infer the structure through their experience and observations. Further, by mentioning specific features, Experiment 1 may have created pragmatic demand characteristics; participants could have been swayed by conversational norms in favor of the narrow scope category (cf., Hilton, 1995). That is, participants might infer a communicative intent in the experimenters' decision to mention only the common feature; had the broad scope category been correct, it would be simpler to communicate that by describing the non-shared feature. The choice to describe the shared feature might indicate that the desired response was narrow-scope.

Accordingly, to provide a stronger test of the narrow latent scope bias in categorization, Experiment 2 investigated a case of visual category learning. Participants first learned about the categories by examining a set of exemplars described visually by a complete set of features. They then made categorization judgments given an exemplar with many different features displayed (rather than described). The design allowed us to observe if the bias persisted in less artificial settings.

Experiment 2

In Experiment 2, participants learned which features were diagnostic of category membership—with categories defined as different species of fictitious monsters—based on a series of images of category exemplars. The study used fictitious creatures rather than existing social categories to ensure that prior knowledge would not influence participant responses, and results would be determined by differences in category structure alone. Participants were shown a set of exemplars from each category to help them learn category boundaries. The narrow scope category had one type of feature in common, while the broad scope category had two types of features in common (see Fig. 1 for an example). After participants had learned about the monster categories by visually inspecting the exemplars, they categorized new exemplars as being a member of one of the categories. On experimental trials, some of the diagnostic information was hidden by an occluding image of a brick wall. Participants had to categorize the exemplar in the absence of this information. The latent scope hypothesis predicted that

participants should prefer to categorize the exemplar as a member of the species of monster that has fewer diagnostic features. To foreshadow the results, the study found strong evidence for narrow latent scope bias.

Method

Participants

Thirty-two participants were recruited through Amazon.com's Mechanical Turk platform, and participated for monetary compensation.

Stimuli

Figs. 1 and 2 provide examples of the category members used in the study.¹ Monsters were created by combining one of five different heads, arms, bodies, and legs. All Bogwomblers had one body part, randomly chosen, in common (e.g., arms in Fig. 1). All Queezlekins had two body parts in common, one identical to the Bogwomblers' defining feature (e.g., arms) and one unique (e.g., legs). Thus, Bogwomblers were considered narrow scope category members since their group membership was determined by only one feature while Queezlekins were considered broad scope category members since their categorization required an additional feature. The remaining body parts were randomly assigned within a counterbalanced design that presented each possible variation twice per creature, for a total of 10 narrow scope creatures (Bogwomblers) and 10 broad scope creatures (Queezlekins), displayed throughout the learning phases of the experiment. The learning phase included two cases where Bogwomblers displayed the non-overlapping feature that was necessary for defining Queezlekin membership. In other words, some examples of the narrow scope category members displayed the features used to define the broad scope category members. Three examples of each creature type were presented side-by-side on an initial introductory page (as displayed in Fig. 1), and the remaining seven creatures were presented individually throughout the learning stage of the experiment.

An additional creature type, labeled as Mugwumps, was included (see Fig. 2). Mugwumps had one feature in common that was distinct from both Bogwomblers and Queezlekins. They were introduced primarily as distractors, and to show participants that the feature common to both the Bogwomblers and Queezlekins was not common to all creatures. For example, Bogwomblers might have all shared the same arms, Queezlekins might have shared the same arms and legs, and Mugwumps might have shared the same torso.

Design and procedure

The experiment consisted of three stages: Instructions, training, and testing. When the study began, participants were told to imagine that they had entered a forest, where they would encounter many strange and dangerous monsters. Their task was to feed any monster that they came across. However, each type of monster would only eat a particular sort of food; participants had to give the creatures the correct food to avoid being eaten themselves. They then received the following critical instructions:

The monsters... have features that can help you figure out which is which. In particular, monsters vary in the kinds of heads, arms, bodies, and legs that they have. While each Bogwombler might look very different from other Bogwomblers, they all have a certain feature or features in common. The same is true for Mugwumps and Queezlekins.

After reading the instructions, participants saw examples of three Mugwumps (see Fig. 2). They were told: "Look at the Mugwumps below. Notice that they all share the same Mugwump [X]," where X was the shared body part (e.g., "torso" in Fig. 2). They were also told that Mugwumps eat Mugwump crackers to facilitate the feeding task

¹ Monsters were drawn by artist Mike LaRiccia in the style of *Black Mane* (LaRiccia, 2005).

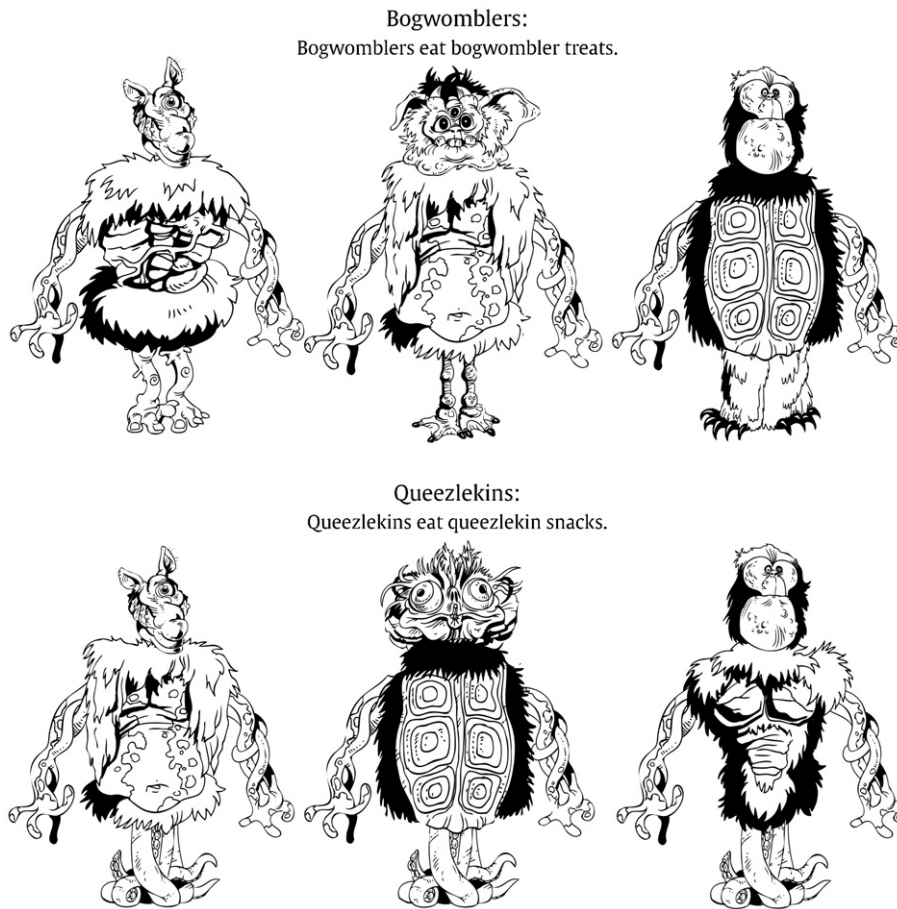


Fig. 1. Examples of Bogwomblers and Queezlekins shown to participants in [Experiment 2](#). Note that Bogwomblers and Queezlekins all share the same arms, while Queezlekins also share the same legs. All remaining body parts have been generated randomly from a set of five possible variations for each.

later in the experiment. This example was included to help participants understand the categorization task, but it used stimuli from the distractor group to avoid interference with learning for the categories being investigated. Participants were required to remain on the example screen for a minimum of 5 s, but could stay as long as they wanted. On the next screen, they saw images of three Bogwomblers and three Queezlekins, along with instructions that read: “Here are some examples of Bogwomblers and Queezlekins. Try to figure out what body parts will help you identify each creature. Remember, you will need to know this information later...” They were also told which type of food

each creature eats (e.g., *Bogwomblers eat Bogwomble treats*). Participants were required to remain on that screen for a minimum of 10 s.

When they chose to proceed, participants entered a training phase, where they were first told which creature was approaching, (e.g., “Look, here comes a Bogwomble”), next they saw the image of a creature, and finally they selected which food to give the creature. They were given feedback on the task and were prompted to correct errors until they selected the correct food (e.g., “The creature won’t eat that. It’s still hungry! What will you feed it now?”). In this phase, participants saw seven Bogwomblers, seven Queezlekins, and two Mugwumps,



Fig. 2. Examples of images of Mugwumps (distractor category members) presented to participants in [Experiment 2](#). In this case, Mugwumps can be identified by their common, trunk-like, torso.

presented in a randomized order. Note that by using stimuli for which participants had no prior knowledge entering the experiment and providing equal numbers of Bogwomblers and Queezelekins within this learning stage, we implicitly equalized the base-rates of each monster within the experimental design.

Finally, participants entered the test phase of the experiment. They read:

... you must continue into the Hidden City on your own. This means that we will not be able to tell you which monsters are which. This might be tricky because the monsters like to hide behind the ruins of the Hidden

City. Luckily, we know that Mugwumps do not live this far into the Lombrozian Forest, so you will only need to carry Bogwomblor treats and Queezelekin snacks...

Participants then saw six monsters—four control and two experimental—presented in random order, and had to choose whether to feed them Bogwomblor treats or Queezelekin snacks in a binary, forced choice task (see Fig. 3 for examples). For control problems, the defining feature unique to the broad scope category (e.g., Queezelekin legs) was visible and either consistent with the broad scope categorization (two problems; broad controls) or inconsistent with the broad scope

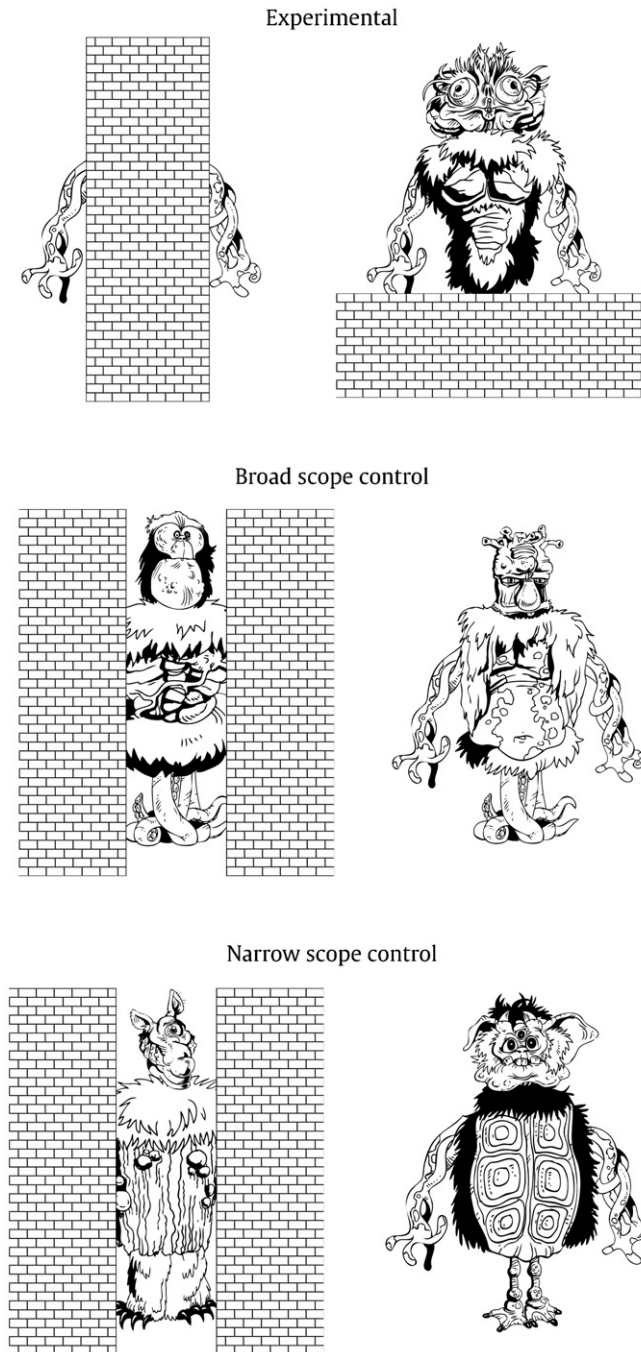


Fig. 3. Examples of image forms used to test categorization of creatures in Experiment 2, given stimuli from Fig. 2. In experimental problems, the feature (in this case, the legs) needed to determine whether the creature was a member of the broad or narrow scope category was hidden. In broad control problems, this feature was visible, and consistent with that held by broad scope category members, thereby suggesting broad scope category membership. In narrow control problems, this feature was visible, but inconsistent with that held by broad scope category members, thereby suggesting narrow scope category membership. After viewing each image, participants were asked what type of food they would give to the creature, with judgments based on perceived category membership.

categorization (two problems; narrow controls). For the two experimental cases, the Queezlekins' defining feature was hidden behind a wall so that participants could not objectively determine whether the monster was more likely to be a Bogwomble or a Queezlekin.

Results and discussion

The dependent variable was which food the participant gave to the monster: Bogwomble treats (narrow) versus Queezlekin snacks (broad). Responses to each type of control and experimental question were averaged across trials for each participant based on the percent of responses choosing food for creatures with a narrow scope categorization. In the broad control condition, participants were more likely to select food for monsters with a broad scope categorization (26.6% Bogwomble treats vs. 73.4% Queezlekin snacks, a one-sample Wilcoxon test relative to 50%, $z = 3.44$, $p = .001$). In the narrow control condition, participants were more likely to select food for monsters with a narrow scope categorization (75.0% Bogwomble treats vs. 25.0% Queezlekin snacks, one-sample Wilcoxon test relative to 50%, $z = 3.02$, $p = .002$). Both of these patterns of response can be construed as normative, and they demonstrate that the task made sense and that participants were paying attention and answering sensibly.

Critically, in response to experimental questions, where the distinguishing feature of the broad scope monster was hidden and there was no objectively correct answer, participants were more likely to select food for monsters with a narrow latent scope categorization (68.8% Bogwomble treats vs. 31.2% Queezlekin snacks; one-sample Wilcoxon test relative to 50%, $z = 2.56$, $p = .011$). This finding is consistent with the hypothesis that when there is a lack of information that can be used to discriminate between a broad and narrow latent scope categorization, people are biased towards choosing the category with narrower latent scope. By relying on participants to learn category boundaries through a series of visual examples, the experiment rules out many pragmatic demands that accompany verbal descriptions of categories, and provides converging evidence for a robust narrow latent scope bias across categorization tasks.

General discussion

Across both verbal and visual categorization tasks, we presented participants with limited, uncertain categorical information and demonstrated that, all else being equal, they were consistently biased to categorize exemplars into categories with narrower latent scope. In a verbal social categorization task, participants were biased towards believing that a novel exemplar belonged to a category with narrower latent scope. The visual categorization task showed that the narrow latent scope bias persisted when participants relied solely on visual cues to determine category boundaries, and when subjects learned the categories rather than being told explicitly about their features. It also ruled out the possibility that pragmatic demands of verbal descriptions lead to the narrow scope bias.

Previous research revealed a bias towards explanations with narrower latent scope (Khemlani et al., 2011), and the present studies show that such a bias occurs in categorization tasks as well. In each case, people responded to uncertainty in their environment by choosing the category or explanation that most closely matched the known information. This preference held despite instructions stating that participants should not infer a cue's absence from uncertainty.

On the surface, this finding may appear to contradict previous findings in the literature. For example, Patalano, Chin-Parker, and Ross (2006) demonstrated that people prefer to use more coherent categories in making predictions about others, and one feature of coherent categories is that they have more features. While this could seem a contradiction to the present predictions, there is a striking methodological difference in the two approaches. In our studies, participants were instructed to choose which category an exemplar was in given

incomplete information about the exemplar's features and unknown category membership. In the Patalano et al. paradigm, an item was known to be in multiple categories, and participants chose which category's features it had. Thus, it appears that lack of knowledge of category membership is a key element for eliciting the narrow scope bias.

The present experiments ruled out several possible mechanistic accounts. For features that were diagnostic solely of the broad scope category, verbal experiments were ambiguous about whether it was possible for the parallel feature to be held by members of the narrow scope category. For example, in the verbal categorization task, participants were told: "Both hunters and spear fishermen carry spears, but spear fishermen also carry nets." Participants could have interpreted this to mean that hunters *do not* carry nets. In this case, the narrow latent scope bias might be related to processing negative information rather than scope, since both the narrow and broad latent scope categories could be defined by two features. If "does not carry a net" serves as a second defining feature for the narrow latent scope category, then the scope of the two categories would be identical, and the category we have referred to as the narrow scope category would instead be differentiated by the presence of a negatively defined feature (i.e., a feature described by its absence). The visual categorization task rules out this possibility in two ways. First, two of the ten examples of the narrow scope category members shown during the learning stages of the experiment displayed the features used to define the broad scope category members. This manipulation clarified that the narrow scope category members *can* possess features that define broad scope category members. Second, for the remaining eight examples, participants saw a range of features for that body part, thus clarifying that there is no specific "opposite" feature defining the narrow latent scope category.

Another concern is that participants may believe that narrow latent scope category members are more common than broad scope category members, and that the inferred base rate is responsible for the bias. However, base rates were equalized in the present studies. **Experiment 1** equalized base rates explicitly and **Experiment 2** did so implicitly through equal representation of narrow and broad scope category members throughout the learning stage. The bias's presence despite equal base rate information is consistent with a large literature demonstrating that people frequently neglect base rates (e.g., Kahneman & Tversky, 1973). Prior research on the scope bias likewise showed that explicit statement of base rates did not affect evaluations of explanations that differed in scope (Khemlani et al., 2011, Experiments 1d and 2).

One concern that readers may have is that the possibility that the observed pattern could in fact conform to normative predictions. The distinguishing feature needed to classify broad scope category members (e.g., Queezlekin legs) was limited to a single feature type for broad scope category members (e.g., tentacles), but was unrestricted for narrow scope members (e.g., tentacles, stumps, claws, muscles). Had reasoners relied on this information alone, they could have surmised that the likelihood of the creature having tentacles rather than one of the other four leg options is small, and thus the likelihood of the unidentified creature being part of the broad scope category is also small. However, the reasoning, while intuitive, requires a misrepresentation of the relevant conditional probabilities related to biases from partition dependence (e.g., Fox & Clemen, 2005; Fox & Rottenstreich, 2003). That is, the logic stems from an implicit partition to the problem at the level of the individual feature rather than at the level of the category. Given the study design and the equal number of appearances of broad and narrow scope members throughout, broad scope members' defining characteristics appeared the same number of times as all of the other possibilities in combination. Thus, the equal likelihood of each type of monster appearing should be the dominant factor in participants' judgments. This factor renders the number of possible features that could satisfy category membership inconsequential.

Why are people biased towards narrower latent scope? One possibility is that participants prefer to categorize objects based on

representativeness (Kahneman & Tversky, 1973). In this case, the narrow scope categories require fewer features to determine category membership, and are thus more representative of the limited data provided. Similarly, participants may favor categories for which larger proportions of the features are observed. In some trials, participants saw evidence for half of the data needed to define a narrow scope category member, but only one-third of the evidence to define a broad scope one. Broad scope categories might be penalized for encompassing more unobserved features than narrow scope explanations. The behavior could be a specific manifestation of cognitive economy, whereby people save resources by looking for the closest match to observed data.

Note once again that this process is not normative: compare the experiment to a situation where one urn (category A) contains 3 red and 3 green balls while another (category B) contains 3 red, 3 green, and 3 yellow balls. Participants seem to be responding as though they learned that a red ball was randomly picked from one of the urns and they must determine from which urn the ball was picked. Instead, the information presented to participants in these experiments parallels a situation in which a red-ball-detector identifies the presence of a red ball, and they must determine which urn the red-ball-detector is more likely to be near. Although in the former case, narrow scope bias would be normatively appropriate (you are more likely to draw a red ball from an urn where 50% are red rather than 33% are red), in this latter case, there is no way to distinguish between the two likelihoods (a detector is equally likely to show the presence of a red ball regardless of the proportion of red balls in the urn).

When evaluating others, people are frequently faced with insufficient information to make category judgments with certainty. Nonetheless, people make inferences about category membership and rely on category-level representations to aid the process of person perception. Experiment 1 and Experiment 2 showed that people preferred to categorize others into groups with narrower latent scope when presented with uncertain information. They suggest that people are most likely to infer that others do not have a particular attribute if that attribute is not directly observed even if that attribute is, in fact, unobservable.

The present findings suggest that the narrow latent scope bias is likely to persist across an array of social and cognitive processes that involve causal reasoning about uncertainty, and they have implications for a wide range of judgment tasks. People frequently are required to make judgments and decisions with only limited information available, whether determining expectations about peers, assessing the treatment of others, providing eyewitness testimony, or making medical diagnoses. A greater understanding of how people respond to limited information can help explain and predict stereotypes and inferences about others.

Acknowledgments

We are grateful to members of the Opplab, Sam Glucksberg, Adele Goldberg, Todd Gureckis, Matt Johnson, Phil Johnson-Laird, Max Lotstein, Greg Murphy, Marco Ragni, Bob Rehder, Joe Simmons, and Laura Suttle for useful feedback. We thank Michael LaRiccia for his help in creating the materials for the experiment and David Mackenzie for his help with programming. This research was partially supported by a National Science Foundation Graduate Research Fellowship awarded to the second author.

Appendix A. Basic problem forms for Experiment 1

This chart provides a basic problem form that is representative of the information that participants received for each of the possible control and experimental problems. A and B denoted two possible categories,

C denoted the person in question, and X, Y, and Z denoted distinct features.

Control premises	Experimental premises
As have X Bs have X and Y C has Y; we don't know whether or not he has X.	As have X Bs have X and Y C has X; we don't know whether or not he has Y.
As have X Bs have X and Y C has X and Y.	As have X and Y Bs have X, Y, and Z C has X; we don't know whether or not he has Y or Z.
As have X and Y Bs have X, Y, and Z C has X, Y, and Z.	As have X and Y Bs have X, Y, and Z C has Y; we don't know whether or not he has X or Z.
As have X and Y Bs have X, Y, and Z C has Z; we don't know whether or not he has X or Y.	As have X and Y Bs have X, Y, and Z C has X and Y; we don't know whether or not he has Z.

Appendix B. Category descriptions used in Experiment 1

1. In the jungles of the Amazon about half of the Tokolo tribe members are hunters, and the other half are spear fishermen. Both hunters and spear fishermen carry spears, but spear fishermen also carry nets.
2. About half of people have gene MNR834 and the other half have gene ANB349. Everyone with gene MNR834 has the protein Nyerlon in their hair and the chemical Mercien in their blood. Everyone with ANB349 has the protein Nyerlon in their hair and the chemical Mercien in their blood, and the enzyme Entlene in their saliva.
3. The James Polk high school marching band has different uniforms for different band members. Trumpet players have a star on the front of their shirt, while trombone players have a star on the front of their shirt, and a moon on the back of their shirt. The band has equal numbers of trumpet and trombone players.
4. Half of students at the annual school geography bee are in Ms. Magaletti's class, and half are in Ms. Anders' class. Ms. Magaletti's class has covered European geography, while Ms. Anders class has covered European geography and Asian geography.
5. At a scrimmage between two high school basketball teams, half the players are from Oakland High and the other half are from Ridgewood High. Accidentally, both teams are wearing red jerseys that have black numbers on them. But, Ridgewood won the championship last year, and players from Ridgewood are also wearing a championship ring.
6. Among the guests at a wedding are family members of both the bride and groom – about equal numbers of each. The bride's family has blonde hair and deep voices, while the groom's family has blonde hair, deep voices, and Scandinavian accents.
7. In the forest of Zaire, about half of the Innuri tribe members (Legetarians) eat furrel meat and pled meat. The other half (Rendetarians) eat furrel meat, pled meat, and herdoz meat.
8. Two tech startups – xomato.com and uxwire.com – share an office building. Each has half of the building and about the same number of employees, but they have different perks. Employees at xomato.com get driven to and from work in a limo everyday. Employees at uxwire.com get driven to and from work in a limo everyday and also get free breakfast once they arrive.

References

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, 41, 361–416.
 Cantor, N., & Mischel, W. (1979). Prototypes in person perception. *Advances in Experimental Social Psychology*, 3, 4–47.

- Cohen, C. E. (1981). Person categories and social perception: Testing some boundaries of the processing effect of prior knowledge. *Journal of Personality and Social Psychology*, 40(3), 441.
- Crisp, R. J., & Hewstone, M. (2007). Multiple social categorization. *Advances in Experimental Social Psychology*, 39, 163–254.
- Fox, C. R., & Clemen, R. T. (2005). Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior. *Management Science*, 51(9), 1417–1432.
- Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, 14(3), 195–200.
- Griffiths, O., Hayes, B. K., & Newell, B. R. (2012). Feature-based versus category-based induction with uncertain categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 576–595.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, 13(2), 141–154.
- Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin*, 118, 248–271.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237–251.
- Khemlani, S. S., Sussman, A.B., & Oppenheimer, D.M. (2011). Harry Potter and the sorcerer's scope: Latent scope biases in explanatory reasoning. *Memory & Cognition*, 39(1).
- Klein, S. B., Loftus, J., Trafton, J. G., & Fuhrman, R. W. (1992). Use of exemplars and abstractions in trait judgments: A model of trait knowledge about the self and others. *Journal of Personality and Social Psychology*, 63(5), 739.
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, 14(4), 551–577.
- LaRiccía, Michael V. (2005). *Black mane*. Suffolk, UK: One Time Press.
- Lombrozo, T. (2009). Explanation and categorization: How "why?" informs "what". *Cognition*, 110, 248–253.
- Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual review of psychology*, 51(1), 93–120.
- Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1995). The dissection of selection in person perception: inhibitory processes in social stereotyping. *Journal of Personality and Social Psychology*, 69(3), 397.
- Malt, B. C., Ross, B. H., & Murphy, G. L. (1995). Predicting features for members of natural categories when categorization is uncertain. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 646–661.
- Molden, D. C., & Higgins, E. T. (2004). Categorization under uncertainty: Resolving vagueness and ambiguity with eager versus vigilant strategies. *Social Cognition*, 22, 248–277.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 289–316.
- Murphy, G. L., & Ross, B. (1994). Predictions from uncertain categorizations. *Cognitive Psychology*, 27, 148–193.
- Murphy, G. L., & Ross, B. (2005). The two faces of typicality in category-based induction. *Cognition*, 95, 175–200.
- Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5, 411–419.
- Patalano, A. L., Chin-Parker, S., & Ross, B. (2006). The importance of being coherent: Category coherence, cross-classification, and reasoning. *Journal of Memory and Language*, 54, 407–424.
- Rehder, B. (2003a). Categorization as causal reasoning. *Cognitive Science*, 27, 709–748.
- Rehder, B. (2003b). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1141–1159.
- Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs of categorization, induction, and similarity. *Journal of Experimental Psychology: General*, 323–360.
- Ross, B. H., & Murphy, G. L. (1996). Category-based predictions: Influence of uncertainty and feature associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 736–753.
- Sloman, S. A., Love, B. C., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 189–228.
- Stangor, C., Lynch, L., Duan, C., & Glass, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology*, 207–218.
- Tajfel, H., Billig, M. G., Bundy, R. P., & Claude, F. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology*, 149–177.
- Verde, M. F., Murphy, G. L., & Ross, B. H. (2005). Influence of multiple categories on the prediction of unknown properties. *Memory and Cognition*, 33, 479–487.
- Waldmann, M. R., Holyoak, K. J., & Fratianne, A. (1995). Causal models and the acquisition of category structure. *Journal of Experimental Psychology: General*, 124(2), 181.