



City Research Online

City, University of London Institutional Repository

Citation: Hampton, J. A., Dubois, D. and Yeh, W. (2006). Effects of classification context on categorization in natural categories. *Memory & Cognition*, 34(7), pp. 1431-1443. doi: 10.3758/BF03195908

This is the unspecified version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <http://openaccess.city.ac.uk/1000/>

Link to published version: <http://dx.doi.org/10.3758/BF03195908>

Copyright and reuse: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

To appear in *Memory & Cognition* 2006

Effects of Classification Context on Categorization in Natural Categories

James A. Hampton,
City University, London

Danièle Dubois,
CNRS, Paris

and

Wenchi Yeh,
Hewlett Packard Company

Memory & Cognition, 2006, in press

Corresponding author:

James A. Hampton

Psychology Department, City University

Northampton Square, London EC1V OHB, UK

Tel: 00 44 20 7040 8520

hampton@city.ac.uk

Effects of Classification Context on Categorization in Natural Categories

The pattern of classification of borderline instances of eight common taxonomic categories was examined under three different instructional conditions to test two predictions. First that lack of a specified context contributes to vagueness in categorization, and second that altering the purpose of classification can lead to greater or lesser dependence on similarity in classification. Instructional conditions contrasted purely pragmatic with more technical/quasi legal contexts as purposes for classification, and these were compared with a no-context control. Measures of category vagueness comprised between-subjects disagreement and within-subjects consistency, while measures of similarity-based categorization comprised category breadth and the correlation of instance categorization probability with mean rated typicality independently measured in a neutral context. Contrary to predictions, none of the measures of vagueness, reliability, category breadth, or correlation with typicality were generally affected by the instructional setting as a function of pragmatic versus technical purposes. Only one sub-condition in which a situational context was implied in addition to a purposive context produced a significant change in categorization. Further experiments demonstrated that the effect of context was not increased through having participants talk their way through the task, and that a technical context did not show any more all-or-none categorization than a pragmatic context. The findings place an important boundary condition on the effects of instructional context on conceptual categorization.

A phenomenon of major importance for psychological theories of concepts is the "vagueness" of many of our conceptual categories. While every category can be said to have clear members (for example a chair is clearly a type of furniture) and clear non-members (a cucumber is clearly not a type of furniture), there are also instances which are borderline to a category. When asked to decide whether rugs, paintings or televisions are types of furniture, people are frequently uncertain about the answer. There is a vagueness in our use of common language terms that arguably makes such questions undecidable. The problem of vagueness poses serious threats to many accounts of the semantics of natural language (Keefe & Smith, 1997; Osherson & Smith, 1997), and so the issue of what gives rise to the phenomenon is of central importance to theories of cognition.

There have been many demonstrations of vagueness. For example McCloskey and Glucksberg (1978) presented two groups of students with lists of words each headed by a category name such as fruit or fish. One group was asked to give typicality ratings, saying how typical or representative each word was of the category as a whole. The other group made a simple Yes/No categorization decision about each word, and returned four weeks later to make the same decision again. Many items in the lists showed a high level of disagreement between participants, and poor test-retest reliability or consistency. These items also tended to be borderline in terms of their rated typicality in the category.

In a subsequent re-analysis of McCloskey and Glucksberg's data, Hampton (1998) showed that categorization probability for an item was closely related to rated typicality by a simple monotonically increasing threshold function. List items that deviated from this standard function tended to be unfamiliar, or might be a part or property of an instance rather than an instance itself. For biological categories cases also could deviate from the function if they had the appearance of a category member without technically belonging to it, or if conversely they technically belonged to a category while not sharing appearance features. From this analysis, Hampton (1998) argued that categorization decisions are to a large extent based on the same "family resemblance" semantic information as is used in judging typicality. An item is judged a category member if the similarity between the item and the prototype for the category passes some threshold value.

Because both the concepts retrieved and the threshold criterion may vary across occasions, the probability of categorization rises as a monotone threshold function of the semantic similarity of the instance and category concepts.

Why should categorization be so unstable at the category borderline? Barsalou (1987) argued that instability could reflect variation between individuals in their conceptual representations or in their recent experience with a category. However it is also possible that instability in categorization results from the lack of a specific context with respect to which the categorization has to be made. In every day language, words are used in specific contexts with specific communicative goals, and this contextual support is missing in standard categorization experiments. If individuals respond to the lack of context by arbitrarily constructing one of their own, then differences in the resulting conceptual representations would create instability.

The purpose of the present study was to investigate variation in categorization resulting from a particular type of contextual source. It has been argued (Braisby, 1993; Braisby & Franks, 1997, 2000; Braisby, Franks & Harris, 1997) that a major source of instability and vagueness in categorization judgments is the lack of any explicit context for the categorization. If asked whether a television is furniture, someone may give a different reply if the question is asked in the context of designing the look of a living room, as opposed to planning the need for electrical outlets in the home. The purpose for which a classification is made may be crucial to how it is performed. Braisby and Franks (1997) went so far as to argue that the lack of a clear context, or perspective, may be a major reason that categories appear to be so vague (see also Rey, 1983). According to their position, the observed vagueness of categories is in large part the result of categorizers selecting at random different well-defined concepts relevant to different contexts or “perspectives”. As individuals recruit their own default context to the task, so differences of opinion about categorization may be more apparent than real. The hypothesis resonates with Barsalou’s (1987) proposal that people construct category representations “on the fly” as different tasks are presented to them, so that there is inherent instability in the information represented in working memory on any one occasion of categorization.

Several studies have deliberately manipulated context in categorization tasks (for a review

see Murphy, 2002, pp. 413-422). For example Roth and Shoben (1983) varied sentential contexts as in "The bird crossed the farmyard", and showed through measures of sentence processing that the context could lead to reversal of normal typicality effects (for example "chicken" was read faster than "robin"). Barsalou and Sewell (Barsalou, 1987) showed that asking participants to take the point of view of (say) a suburban housewife as opposed to a "red-neck" farmer produced marked changes in the typicality ranking of instances within categories such as vehicles or foods. In both these studies, situational context was manipulated, and marked effects were observed on the relative typicality and ease of processing of different category instances. Another important demonstration of shifts in categorization with context was provided by Medin et al. (1997). Different groups of tree experts sorted trees by similarity. Depending on whether the experts were taxonomists, landscape gardeners or parks maintenance staff, the structure observed was very different. Each group had important dimensions of similarity not used by the others that were relevant to their own profession.

Our experiments differed from these studies in several ways. First, we focused not on the situational context of a classification, but rather on its purpose. It is clear that when considering animals in the situational context of eating them as opposed to animals in the context of inviting them into your home, you will adopt very different views of what instances make typical candidates. However it is a largely unexplored question whether, differences in expertise aside, changing the purpose for which a classification is to be used will generate such shifts. Second, our experiments differed from many previous studies in that rather than measuring typicality structure, we measured changes in categorization itself. We wished to explore whether an item would be considered to be a member of the category in one purposive context but not a member of the same category in another. If category vagueness is in part owing to contextual ambiguity, it is clearly important to show that context can affect categorization decisions as well as typicality structures. Finally we wished to explore the possibility that disagreement and inconsistency would be reduced when a clear purposive context for classification was provided.

In an unpublished study, Braisby & Franks (2000) found evidence that categorization could be strongly influenced by shifting perspectives. They contrasted two types of borderline

instance for natural kind classes – those that had the appearance but not the essence (for example an Easter egg as an egg), and those that had the essence but not the appearance (for example a scrambled egg as an egg.) People were asked to judge whether it was appropriate to use the word “egg” to talk about each object. In a series of studies they showed that the relative frequency with which these two types of instance were categorized as eggs depended on a number of factors. Taking the perspective of a sculptor rather than a biologist naturally put more weight on appearance. Imagining speaking to an adult non-native speaker led to more weight for essence, whereas imagining speaking to a 4-year-old child led to more weight on appearance. If the purpose of using the word was in conversation or in defining a meaning more weight was placed on essence, whereas using a word for picking an object from an array switched the weight to appearance. Focusing on “true classification” rather than appropriateness of word use shifted weight onto essence. It is clear therefore that communicative setting and purpose can be very influential in affecting categorization, at least as shown in people’s judgments of appropriate words to use for objects. Our aim was to determine how general this effect may be, looking at a wider range of categories and borderline cases, and additionally measuring the consistency with which people make their judgments.

To test whether vagueness is owing to a lack of information about the purpose of classification, we aimed to provide participants with a clear perspective from which to make their categorizations. If the purpose of categorization is made clear, then there should be less vagueness. The first experiment therefore employed three categorization conditions. One condition (the No-Context Control) was a simple context-free categorization task, while the remaining two conditions offered different scenarios, explaining the purpose and importance of the classifications. The first prediction for this factor was that providing a specific context of some kind would lead to less individual disagreement and inconsistency, that is, less vagueness.

The second prediction concerned the kind of categorization context provided, and the degree to which categorization would be dissociable from similarity to the category prototype. Hampton (1998) argued that the degree to which categorization probability is a simple monotone function of typicality can be taken as a test of the degree to which participants are simply

categorizing on the basis of similarity to prototype, as opposed to using a more complex explanation-based or theory-driven decision process. For example in biological categories (but not others) categorization probability was not well correlated with typicality for items that were technically in categories different from their appearance (for example whales and bats).

The importance of causal-explanatory theories for categorization has been well established (Ahn et al., 2000, Murphy & Medin, 1985). Rips (1989) has argued that similarity is only a crude approximation to the basis on which people categorize the world. Under the right circumstances it is possible to show that categorization and typicality judgments may be dissociated, as people turn for categorization to deeper core information about a concept and ignore superficial appearance (Ahn & Dennis, 2001). We hypothesized that depending on the classification context, similarity may turn out to be a more or less appropriate basis for categorization. For example when setting up a news/interest group on the subject of fish, it would be appropriate to use a “loose” category of fish, that could include shell fish or dolphins along with “true” fish such as cod or trout. On the other hand when preparing a scientific report on the ecological status of different species, it would be more appropriate to use a quasi-biological definition for fish, which would exclude these items, but might instead include seahorses. Our choice of contexts was designed to take advantage of this intuition that there may be more technical and more pragmatic forms of categorization. In the Technical Condition, the purpose of categorization had a technical foundation, based on scientific or legislative goals, while in the Pragmatic Condition, the purpose was more loosely practical, based on providing a classification that would be easy to use, and match people’s general expectations. We predicted that a pragmatic context would be more likely to reveal similarity-based categorization, since it would place things in categories where the mass of people would expect to find them. Categorization would therefore rely more on the “identification schema” or prototype of a concept, and less on its core definition. A technical context however should encourage participants to use deeper causal-explanatory schemas for categorization, in which the relation of categorization to typicality was less direct. In order to test the degree to which categorization was based on similarity, correlations were calculated across items of the probability of categorization with the mean rated typicality provided by an independent group of participants.

A further measure to be compared between these conditions was the overall threshold criterion used for categorization in each category (Hampton, 1995). We expected that in a pragmatic context, people would take a broad view of what may be included in a category, whereas in a technical context, the category boundary would be drawn more tightly. The No-Context condition was predicted to be intermediate between these two.

A final source of interest in the task came from possible differences between categories in the degree to which they would be affected by contextual instructions. There has been considerable interest in domain differences in the way that concepts are represented (Barr & Caplan, 1987; Estes, 2003; Kalish, 1995). Although the design of the experiment was too small to permit adequate sampling from different semantic domains, we deliberately chose categories from four different ontological domains in order to provide a broad range of materials. Two biological kinds, fish and insects, were expected to show marked differences between technical and other contexts, because of the existence of biological definitions for these terms. Two artifact kinds, tools and furniture were expected also to show changes across condition on the basis that technical contexts would place greater weight on the utility or function of the objects, which may provide the central core of artifact concepts (Bloom, 1996). Finally there were two categories of edible plants, fruits and vegetables, and two social activities, sports and sciences, where the technical contexts were expected to tap other possible forms of theoretical knowledge and beliefs. While these pairs of semantic categories were clearly not unbiased samples of their respective domains, it was intended that analysis at the level of individual categories could provide indicative evidence of any strong and systematic domain differences that may exist.

To recap, we predicted first that adding any context at all would reduce vagueness, and second that technical contexts would contrast with pragmatic contrasts by inducing tighter category boundaries and a reduction in the dependence of categorization on similarity, as seen in the correlation with typicality.

Experiment 1

Method

Participants. 100 students at the University of Chicago volunteered for the experiment in

return for a small payment. All were fluent speakers of English.

Materials. The categories and items used are shown in Appendix A. There were 24 items in each of 8 categories. Items were selected from category norms (Battig & Montague, 1969; Hampton & Gardiner, 1983; McCloskey & Glucksberg, 1978; Rosch & Mervis, 1975) and were designed to include clear members and nonmembers of the category, together with a substantial number of possible borderline cases to provide a measure of how vagueness and category membership changes with context. Examples of scenarios for each condition are given in Appendix B. The aim for the Technical Condition was to provide a legalistic or scientific context stressing the important consequences of making a correct classification. For fish and insects the classification was to be used by a government agency for monitoring the ecological performance of different nations. For fruits and vegetables, the classification was concerned with economics and trade. For furniture and tools the context involved tax regulations, while for sports and science the context was concerned with appropriate use of funds by government agencies. In the Pragmatic Condition, the stories were concerned with placing things in categories where people would expect to find them, so that they would be easily found. A variety of contexts were used, including an Internet news group (for fish and insects), a mail-order catalogue (for fruits and vegetables), a department store database for monitoring stock (for furniture and tools) and a library index (for sports and science). Finally, the No-Context condition had the same instructions for all categories: "Consider each of the following items and decide whether they belong in the category of _____".

Design. There were four groups of participants. One group of 40 students provided typicality ratings. Three other groups of 20 students each made Yes/No categorization judgments, according to the three conditions. Categorization was retested 3-4 weeks later.

Procedure. Participants were given booklets to complete under supervision. Order of categories within booklets and words within categories was balanced. Participants were asked to read the scenario and judge each item by circling a choice of Y (Yes), N (No) or Ø (meaning of word unknown). After 3 to 4 weeks, participants repeated the task with instructions to make fresh judgments without trying to recall earlier decisions. A new random order of categories was used.

The Typicality Rating condition used standard typicality instructions, using a scale from 1 to 10 (10 most typical) following McCloskey & Glucksberg (1978):

"...rate each word according to how typical or atypical it is as a member of the category
...decide how good or representative an example each word is of the category named..."

An example was then given for the category of flowers. Item and category order were counterbalanced as for the other conditions. Typicality ratings were not subjected to a retest.

Results

For each item in each of the 3 categorization conditions, the probability of a yes response was calculated, based on the two responses made by each participant in each condition. Reliability of the probabilities (mean .99) and the mean typicality ratings (mean .96) was uniformly high.

Inter-Subject Agreement. The first measure of category vagueness was inter-subject agreement. The addition of contextual instructions was predicted to reduce contextual ambiguity so that people would agree more on the categorization of items. The left-hand panel of Table 1 shows disagreement measured as the proportion of Non-Modal Responses, NMR, the proportion of participants giving a No response when the majority said Yes, or saying Yes when the majority said No. Overall, as predicted there was slightly more disagreement in the No-Context condition (18.6%) than in the other conditions (17.1% and 18.0%, SE 1%). These means are for all items, including clear members and non-members. NMR rose to a maximum of around 35% for items at the center of the typicality scale (as in McCloskey & Glucksberg, 1978), leaving plenty of opportunity for a reduction in its value. ANOVAs by subjects and items were run with Category and Condition as factors. There was no overall effect of condition ($F_1 < 1$), and nor was the planned comparison between the No Context Condition and the others significant ($t(57) = 0.9$, NS). There was therefore no evidence that providing a context increased inter-subject agreement. With $\alpha = .05$, estimated power to detect a difference in condition means of as much as 5% between the No-Context and the other conditions was greater than 97%. (In reporting power estimations, the lower of the two powers (items or subjects) is always quoted.)

Otherwise, there was a significant main effect of category, ($F_1(7,399) = 9.78$, $p < .001$, $F_2(7,184) = 2.34$, $p = .026$), attributable to greater agreement about the four biological categories

than about the activity and artifact categories (mean NMR = 14.6 and 21.2 respectively). There was no interaction between condition and category. In summary, in none of the categories was the No-Context condition clearly subject to greater disagreement than all the others. Our expectation that a more specific categorization context would generally reduce category vagueness as indexed by levels of disagreement was not supported.

INSERT TABLE 1 ABOUT HERE

Within Subject Consistency. The second measure of vagueness was the test-retest consistency of categorization responses. Addition of a context was predicted to improve the stability of categorization decisions by reducing contextual ambiguity. If an item is categorized in relation to a contextualized concept, then recall of the same context should facilitate consistent responding. Note that if different individuals contextualized the concepts in idiosyncratic ways, then NMR would remain high, but consistency should still improve relative to the No-Context condition. The right-hand panel of Table 1 shows consistency as the percentage of responses that were the same on retest. On average, the same response was given 90% of the time. This level of consistency compares with a mean level of 88% in McCloskey and Glucksberg's (1978) data, and includes clear members and non-members. At the middle of the typicality scale mean consistency fell to around 82%. There was no tendency for contexts to increase consistency. Mean values across conditions varied very little, from 91% for the Pragmatic condition to 89% for the Technical condition (SE = 0.7%). In ANOVAs, only the main effect of category was significant in both analyses. Estimated power to detect the contrast between the No-Context and the other conditions was 70% for a difference in means of 3%, and 97% for a difference of 5% (alpha = .05, 2-tailed). The effect of category again showed up as a difference between the biological and food categories (92%) and the others (88%). In summary, the main prediction of lower consistency in the No-Context control condition was not supported.

Relation of Categorization to Typicality. The next analysis considered the correlation between categorization probability in each of the categorization contexts and the mean rated typicality of the items. To the extent that a high correlation is observed between categorization probability and typicality, it may be concluded that categorization is based on similarity, and for

this reason we had argued that Pragmatic Contexts should show higher correlations with typicality than Technical Contexts. To illustrate these correlations for “default” categorization, Figure 1 shows scatter plots for each of the categories between probability of categorization in the No-Context condition (y-axis) and rated typicality (x-axis). It can be seen that for all categories except for sport, there was a strong and systematic relation between the two measures. The left-hand panel of Table 2 shows the correlations between typicality and categorization probability. All correlations were high (the mean correlation of .95 was close to the theoretical maximum imposed by the reliabilities of the measures), and differences between the three overall correlations by condition were slight, and not statistically significant ($\chi^2(2) = 2.01, p > .10$, Weatherburn, 1961, pp. 203-205). Estimated power to detect a difference in mean correlations significant at .05 between any two conditions of as little as .95 versus .90 was over 90%. Taking the full set of 24 correlations as a whole however, there was evidence for non-homogeneity ($\chi^2(23) = 43.4, p < .01$), suggesting that some individual correlations were significantly lower than the rest.

INSERT FIGURE 1 AND TABLE 2 ABOUT HERE

In this instance, it was quite easy to find plausible accounts for the cells in the correlation matrix with lower coefficients. The biological and food categories had uniformly high correlations with typicality, (.94 or greater). Categorization of sports correlated less well with Typicality except in the Pragmatic Condition, whereas for tools and furniture it was the Pragmatic Condition that showed a lower correlation. Borderline sports activities could be divided into two groups, those involving physical exercise but little skill (for example aerobics and jogging), and those involving skill but no physical effort (croquet, billiards, and darts). In the Pragmatic condition rated typicality was a good predictor of categorization probability for both groups. For the other contexts however, the categorization placed more emphasis on skills than on physical effort, so that rated typicality over-predicted the categorization of aerobics and jogging, and under-predicted the categorization of croquet, billiards, and darts as sports. The upper panel in Figure 2 shows the effect for the Technical Condition, in which a categorization was required for use in Sports Funding decisions. The similarity between the Technical and the No-Context

conditions for sports suggests interestingly that the default categorization of sports also places greater weight on the skills aspects of sports than is seen in typicality judgments.

INSERT FIGURE 2 ABOUT HERE

Tools and furniture showed reduced correlations with typicality in the Pragmatic Condition only. These scenarios involved devising a database for employees in a department store, so that they could check on the availability of different items. Participants appear to have adopted the layout of department stores as a guide to categorization. For example, electrical appliances of various kinds were less likely to be counted as furniture, since in many department stores they would be in the electrical goods department rather than the furniture department (see lower panel, Figure 2). This strategy was also accompanied by higher consistency for these two categories in the Pragmatic Condition (.90 and .93) compared with the other conditions (values ranging from .85 to .89), and for tools it also was accompanied by a reduction in Non Modal Responses. Note that the reduction in vagueness and the reduced correlation with typicality were produced only in those context-category combinations that allowed a situational context to be imagined. Providing a purpose for the classification per se had no discernible effect overall, yet the dependent measures were sensitive to situational context effects.

The high overall level of the correlations was obviously affected by the presence in the lists of clear members and non-members. Removing items with less than 10% NMR left between 9 and 19 borderline items in each list. With reduced range and reliability, mean correlation between typicality and categorization fell to 0.85, and the four correlations identified before were the only ones to fall below 0.8.

Criterion. A further important way in which context could affect categorization is in terms of the breadth of the categories. We had predicted that technical contexts might lead to narrower category criteria, as the instructions stressed the importance of producing a fair categorization. The right-hand panel of Table 2 shows the mean proportion of positive categorizations. Contrary to expectation, the No-Context Condition had the tightest criterion overall, with 47% positive responses, compared to 49% in the other conditions (SE 2.6%). In no category did the Technical Condition, which had been predicted to be the tightest, have the lowest number of positive

responses. The probability data (with an arcsine transformation) were subjected to ANOVAs by subjects and items. None of the effects were consistently significant across both analyses. Estimated power to detect a contrast between the No-Context and the other conditions was 70% for a difference of as much as .05 in mean probability, and over 90% for a difference greater than .07 ($df = 57$, $\alpha = .05$).

Discussion

The experiment set out to examine a number of different measures to test the potential effect of a purposive scenario on categorization. First we looked at levels of agreement and consistency in categorization. No systematic effects were found of supplying a contextual scenario. Second, we examined correlations with typicality and category criterion as a test of the idea that "loose" pragmatic contexts would encourage use of default similarity to a common prototype and a broad criterion, whereas more technical contexts would encourage explanation-based categorization and a narrow criterion. The results also failed to support this notion. In fact for 5 of the 8 categories, technical contexts were more similarity-based than were pragmatic contexts.

Only for one category and one context – tools in a department store scenario – did the provision of the context systematically affect all of our measures in the predicted way, by (1) reducing the vagueness in categorization in terms of less disagreement (17% vs 21%) and greater consistency (90% vs. 87%), (2) reducing the correlation with typicality (.86 vs .95) and (3) tightening up the category (45% vs 51%). This consistent pattern is evidence that the manipulation of context can work in the predicted way. The interesting point to note is that the department store context was in many ways very similar to the types of context used by Roth and Shoben (1983), in which provision of a situational setting, such as a farmyard, affected the typicality of different birds. Thus, although we had stressed the purpose of the classification in our instructions, the existence of a familiar situational context was probably the key factor in changing categorization. It remains to be seen whether this one case can be generalized.

Our failure to obtain the predicted effects of context in general across the range of measures, suggests either that participants were ignoring the instructions to imagine themselves in

the given scenarios, or else that their "default" way of thinking of each category was sufficiently powerful to be recruited into the different contexts relatively unchanged. Against the conclusion that the context had no effect at all, one can point to systematic effects observed for sports, furniture and tools. These effects were associated with changes in the categorization of particular subcategories of items, such as exercise sports versus skilled sports, or electrical appliances versus more decorative furniture. The most remarkable finding was that the No-Context condition showed no systematic differences overall from the other conditions in either its correlation with typicality, its category breadth, its between-subject disagreement, or its within-subject consistency. The claim that vagueness in the standard categorization task might generally be the result of contextual ambiguity is hard to reconcile with this demonstration.

One might still argue that the context instruction was not taken sufficiently seriously by the participants. Perhaps they were uninvolved in the task, and so relied on default categorization. Accordingly, in Experiment 2 participants were instructed to speak aloud as they read the context stories and then to spend a minute describing the basis on which they would categorize, before starting to categorize each list. It would be hard under these circumstances for the participants to ignore the stories. Smith and Sloman (1994) for example found that when asked to "think aloud" in one of Rips' (1989) categorization tasks, participants were more likely to show deeper rule-based as opposed to similarity-based reasoning.

Experiment 2

Experiment 2 was a partial replication of Experiment 1. We felt that if an effect of context were to be observed, the most powerful manipulation would be the contrast between the Pragmatic Context, where items should be categorized where people would expect to find them, and the Technical Context, where equitable rules and regulations relating to financial and professional interests were required. These were also the only two conditions where there was any kind of story provided for participants to read out loud. We therefore considered just these two conditions.

Method

Participants. Forty undergraduate students from the same population at the University of

Chicago were paid to participate in the study. None had taken part in Experiment 1.

Design and Materials. The materials and subject population were identical to those of Experiment 1, so comparison could be made between experiments. Twenty participants served in each of the two conditions, Pragmatic and Technical. Each task was only performed once.

Procedure. Participants were given the booklet and asked to explain out loud the situation described. They then spent a minute reflecting on how they were going to approach each task, and what aspects of the category would be important. An example was given of classifying Weapons either for a museum display or for legislating about the legal age for possessing them. Participants then gave Yes/No categorization judgments to the 24 words listed in one of two orders as in Experiment 1. Sessions were tape-recorded.

Results

Inter-subject Agreement. The left-hand panel of Table 3 shows the proportion of NMR. ANOVAs were run by subjects and items with factors of Condition and Category. Only the effect of Condition was significant across both analyses ($F_1(1,184) = 21.08, p < .001, F_2(1,38) = 8.65, p < .01$). NMR was higher in the Technical Condition (21%) than in the Pragmatic Condition (17%, SE 0.3%). As in Experiment 1 the biological categories (17%) showed less disagreement than the others (22%), but in this case the Category factor was not significant by items. Unlike the present experiment, Experiment 1 showed no statistically reliable difference in NMR between the Pragmatic (17%) and Technical (18%) Conditions. An ANOVA with factors of Experiment, Condition and Category showed a two-way interaction between Experiment and Condition that was significant by items ($F_2(1,184) = 9.72, p < .01$), and marginally significant by participants ($F_1(1,114) = 2.9, p < .10$). Comparing Experiments 1 and 2, the level of NMR for the Pragmatic Condition was the same (17%), whereas NMR for the Technical Condition increased from 18% to 21%. Thus one effect of requiring participants to pay greater attention to the context was paradoxically to lead to greater disagreement amongst the participants, but only in the Technical Context condition.

INSERT TABLE 3 ABOUT HERE

Correlation with Typicality. The central panel of Table 3 shows that correlations of

categorization probability with typicality were generally high (mean $r = .95$). If a Technical context leads to more theoretically based categorization, and less emphasis being placed on superficial similarity, then correlations with typicality should be lower for the Technical contexts. In the event, all except one of the categories showed higher correlations with typicality in the Technical Condition than the Pragmatic Condition ($t(7) = 4.29$, $p < .005$, across categories), a trend that was present, but not significant in Experiment 1. Thus the effect of paying more careful attention to the task, and verbalizing the process of performing the categorization was that participants in the Technical Condition appeared to adhere even more closely to default similarity to prototype as the basis for their categorizations.

Examination of individual category data showed that tools and furniture were subject to the same "department store" effect as before, with reduced correlations of categorization with typicality in the Pragmatic condition (.86 and .90, compared with .92 to .97 for the rest). For example the items telephone, dishwasher, refrigerator, and piano were all less likely to be categorized as furniture than predicted by their typicality, and these were all items that would not be normally found in the furniture section of a store.

Criterion. The overall proportion of positive categorizations in each condition was identical (.49) and no different from Experiment 1. The right-hand panel of Table 3 shows the mean categorization probabilities. Categorization probabilities were also compared (a) to the No-Context condition of Experiment 1, and (b) to the conditions in Experiment 1 with the identical context stories. There were no systematic changes discernible, although some categories grew larger and some smaller when participants were required to verbalize the task.

Transcripts. Transcripts of the sessions indicated that participants were clearly aware of the requirements of the task, and had fully understood the scenarios. When asked about the intended basis for categorization, participants in both conditions tended to say that it would be based on the characteristics of the item and their own "gut feeling" about the category membership. There were some attempts in the Technical scenarios to find defining features, but they were not applied systematically. In sum, the transcripts revealed that participants were taking the task seriously and attempting to engage with the scenarios appropriately. However the

categorization data suggest that the basis on which they were categorizing was not markedly different from that which they would use in the absence of any particular context.

Discussion

Experiment 2 replicated two conditions of Experiment 1 under instructional conditions which would encourage greater attention to the contextual manipulation. If the failure of Experiment 1 to show effects of scenario contexts was owing to participants not reading the stories and just proceeding to categorize by default, then we expected the effects originally predicted to emerge in this experiment. The results showed that if anything participants were more likely to rely on their default prototype representation of the categories when the task was made more explicit and verbal. Especially in the Technical Condition participants appeared to adhere even more closely to default similarity to prototype as the basis for their categorizations, and there was an increase in disagreement for this condition compared to Experiment 1. There was no evidence for contextual ambiguity as a source of vague categories. Nor was there evidence for more explanation-based categorization in the Technical Condition, unless the increased level of disagreement in that condition reflected a greater diversity of individual rules being used by the participants. In that case, the categorization probabilities could have resembled the default through the process of averaging across individual differences.

Experiment 3

The third experiment used a different measure of the vagueness of categories. Kalish (1995) presented participants with a categorization scale in which they could either make a clear-cut Yes or No decision, or else make a graded category membership judgment. This procedure was adapted in Experiment 3 to measure whether participants' view of the gradedness of categorization might be influenced by the different contexts. Specifically we expected that in the pragmatic scenarios, where looser similarity-based classification might be considered appropriate, participants would tend to select graded category membership, whereas in the technical scenarios, which emphasized the importance of a correct and fair classification, responses would indicate clear-cut categorization.

Method

Participants. Forty undergraduate students at the University of Chicago were paid a small amount to participate in the study, with 20 in each of the conditions.

Materials. The same categories items and contexts were used as in Experiment 2.

Design and Procedure. Participants were given the same booklets as in Experiment 2 but with a different response scale. Specific instructions (adapted from Kalish, 1995, Study 1) were given on how to use the response scale which consisted of 9 boxes. At one end was a box marked “not at all” and at the other end a box marked “completely”. These response boxes were to be used if an item were clearly in or out of the category, and if category membership was felt to be an all-or-none affair just those options were to be used. Alternatively if a categorization was felt to be a matter of degree, a graded response was to be chosen from the boxes labeled 1 to 7 indicating increasing degrees of membership from “barely”, through “sort of” to “very much”.

Results

Use of Extreme Responses. The prediction of the experiment concerned the use of the two extreme response boxes (“not at all” and “completely”) as opposed to selection of a graded response. Figure 3 shows the proportional use of extreme responses as a function of the typicality of items (taken from Experiment 1) for each of the conditions. Judgments of graded membership were most common in the center of the typicality scale, where they occurred 25-30% of the time. Table 4 shows the percentage of extreme responses given in each condition. The greater overall use of extreme responses for the Pragmatic Condition (85%, S.E. = 2.9%) than for the Technical Condition (81%, S.E. = 2.7%) was not significant in the ANOVA by participants ($F_1 < 1$), nor was the interaction with Category. Estimated power for detecting a difference of 10% or greater between condition means was 80%. The only reliable effect was a main effect of Category ($F_1(7,266) = 4.83, p < .001, F_2(7,184) = 5.18, p < .001$). Extreme responses were more common for biological categories (87%) than for the others (79%) – a result consistent with earlier research (Estes 2003; Kalish, 1995). Table 4 also shows the number of participants who gave extreme responses to all the items in a category. Five of the 40 participants never used graded responses, 3 in the Pragmatic and 2 in the Technical condition. In summary, there was no evidence that a technical context would lead to the perception that the categorizations were less graded.

INSERT FIGURE 3 AND TABLE 4 ABOUT HERE

Discussion

Once again our predictions were unsupported. When participants were adopting a technical context for classification, they were, if anything, less likely to treat the categorization of individual items as all-or-none than when giving a pragmatic classification. If technical contexts promote categorization based on deeper theoretical considerations, then this result is very difficult to explain. Choice of a graded categorization response is generally interpreted as indicating similarity-based classification in which items may fit into a class more or less well (Estes, 2003; Kalish, 1995). It is striking that there was no significant difference in the use of graded responses between the Pragmatic scenarios (where it would be reasonable to use graded categorization) and the Technical scenarios (where the use of the categorization for trade or commercial regulations suggests an all-or-none categorization needs to be used).

General Discussion

The results of the set of three studies described here provide no support for the hypothesis that the vagueness of everyday categorization reflects contextual ambiguity in the purpose for which the classification is being made, and no support for the contention that people would switch to a deeper causal-explanatory basis for categorization when asked to consider categories from a more technical perspective. On the contrary, it appears that similarity, as indexed by context-free typicality judgments, provides a powerful predictor of categorization probability across the groups in all three of the categorization conditions used. Asking participants to pay more explicit attention to the scenarios (Experiment 2) had an effect on the results, but one that certainly showed no sign of bringing them into line with the expected effects. Indeed, the amount of disagreement between participants actually increased and the correlation with typicality was stronger in the Technical Condition in Experiment 2, compared with Experiment 1. Measuring the degree to which people made all-or-none as opposed to graded categorizations (Experiment 3) likewise showed no evidence that they were more likely to treat technical contexts as all-or-none classification tasks, compared with pragmatic similarity-based classification. In fact throughout the three studies, it was the technical contexts that tended to show the closest relation between

categorization probability and context-free typicality.

What conclusions may be drawn from these studies? First, the notion of a robust "default" conceptual representation for the semantic categories used here seems to gain considerable support. The pattern of categorization probability changed relatively little as a function of different categorization scenarios, suggesting that people were using a similar concept representation in each case. There were some notable effects of condition for particular categories. In Experiments 1 and 2, the Pragmatic Condition which used a department store scenario for tools and furniture generated a different pattern of data for those categories. Even though the classification task referred to creating a stock index, rather than the actual lay-out of items in the store, the existence of this prior system of classification proved too tempting to ignore, and participants tended to classify tools and furniture in terms of what would be found in the corresponding departments within the store. Roth and Shoben (1983) demonstrated a similar effect of situational context on category structure. However whereas they showed that typicality within a category shifted with context, our department store effect was reflected in yes/no categorization of borderline items.

Context also had a systematic effect on the classification of sports, such that when categorizing for the purposes of a library index typicality was highly predictive of categorization, but in other contextual conditions more weight was placed on skill and less on physical effort in classifying activities as Sports. This result however was not apparently owing to ambiguity in the meaning of the concept "Sport." If the context had provided disambiguation, then there should have been less disagreement and greater consistency compared with the No-Context control condition. But this was not the case. It could perhaps be argued that the No-Context condition allows participants to recruit the same default context, whereas the scenarios may have themselves been ambiguous, thus leading to greater disagreement and inconsistency in the context conditions. While this is of course possible, it does not sit easily with explaining vagueness in the absence of context in terms of contextual ambiguity. If the account is unable to predict when contexts will increase vagueness and when they will decrease vagueness, then the explanation is empty.

The overall pattern of our results is clearly one of a failure to find evidence to support our

hypotheses. As such the reader may feel that we have done little but fail to reject the null hypothesis. It is worth therefore pointing out the great contrast between the relatively small and subtle effects of the manipulations of purposive context attempted here and the large and robust effects of other manipulations of context. When concepts are placed in a situational context - birds on the seashore or vehicles seen from a farmer's point of view - then a major restructuring of the typicality of category members takes place (Roth & Shoben, 1983; Barsalou, 1987). Barsalou found very low correlations between the typicality ordering of category items when very different points of view were adopted. Likewise, Medin et al.'s (1997) tree experts generated completely orthogonal structures for classifying of trees depending on the domain of their expertise. We believe that similarly large effects would be observed in our experiments if the scenarios had highlighted one particular subset of items over another. For example if the task had been to consider what activities should count as sport, in the context of a foundation whose aim is to promote public fitness and health, then it is easy to imagine that activities that meet this need (for example disco dancing) would be more likely to be included, and those that do not (for example chess) would be more likely to be excluded. However this was not our aim.

Our primary aim was not only to demonstrate context effects per se, but to test a hypothesis concerning category vagueness - namely that at least a part of the disagreement and instability observed in categorization is owing to a failure to provide a context for the classification, in terms of its purpose or function. We approached the hypothesis by designing scenarios that, taken at face value, presented participants with very different ways of conceiving of the purpose of the categorization. While deliberately not providing a strong bias towards any particular subset of the category or any particular feature of category members, we made the manipulation of the difference between our two types of scenario as strong as we could. For example there is a very clear difference between deciding whether books on economics should be placed in a library index under the general category of Science, and deciding whether a national funding body with responsibility for the support of science should be giving grants for research in economics. One purpose we felt, *prima facie*, called for a pragmatic, similarity-based approach, while the other called for the construction of a classification rule that would need to be defended

and justified.

In our second experiment we took steps to ensure that the scenarios were taken seriously by requiring participants to talk about the task for a minute before starting their classification. In the third experiment we tried a different dependent variable to measure the degree to which participants saw the classification as clear-cut or vague. In none of our studies did we find a lack of purposive context contributing to vagueness in categorization in the standard context-free task as observed in countless previous experiments in this field. The scenarios were direct and easy to follow, but they had no effect on consistency. Nor were the studies underpowered. The fact that subtle effects such as the Pragmatic context effect on Tools and Furniture were identified indicates the power of the experiment to detect effects. Indeed there were significant effects of context (see Experiment 2) but where these occurred they were not in the predicted direction. The results were consistent across categories and experiments, and the high correlations observed between typicality and categorization probability are another indicator of the low error variance in these data. Our measures were accurate, and they revealed that instructional context has little detectable effect on categorization probability.

Taken as a whole, our results strongly suggest that there is a common default way of representing conceptual categories and of making category decisions. No matter whether the classification was being created for a tax regulation, or for a news-group search index, the same underlying pattern of categorization probabilities emerged, and the same degree of vagueness and instability was observed at the category boundary. Context effects in categorization can be readily demonstrated in paradigms where the context invokes a situation with which particular subsets of the category are strongly associated (Barsalou, 1987; Roth & Shoben, 1983). Birds in a farmyard have a different graded typicality structure from birds on the seashore. But they are all still birds. Effects on yes/no categorization are harder to demonstrate.

One of the only successful demonstrations of such effects is the research by Braisby and Franks (2000) described in the Introduction. They deliberately created borderline cases that either shared appearance only with a category (a plastic flower or an Easter egg) or shared essence only (a dried flower or a scrambled egg). The weight given to surface appearance as opposed to

underlying essence in categorizing instances was shown to change as a function of a number of factors affecting the perspective of word use. Our results suggest that there are limits to the generality of these effects. First, vagueness in categorization was not reduced by providing purposive contexts. Second, there was no shift away from similarity-based categorization when more serious consequences could arise from the classification. It is possible therefore that the effects described by Braisby and Franks (2000) are primarily concerned with word use – when should you use the word “egg” to refer to a chocolate egg, or the word “chicken” to refer to a rubber chicken?

We see our results as contributing to the general debate about the stability as opposed to context-dependence of conceptual representations (Barsalou, 1987). At least in respect of changing purposive contexts, concepts appear to be remarkably stable.

Author Notes

This research was made possible by financial support to the first author from the French Ministry for Higher Education and Science, the Nuffield Foundation, the Fullbright Foundation, and the British Academy. Their assistance is gratefully acknowledged. The research was conducted while the first author was on sabbatical leave at the University of Chicago, whose hospitality is also acknowledged. The authors wish to thank Lawrence Barsalou, Nick Braisby, Bradley Franks, Dedre Gentner, Douglas Medin, Gregory Murphy, Lance Rips and Dan Sperber for helpful discussion, and anonymous reviewers for helpful suggestions.

References

- Ahn, W. & Dennis, M. J. (2001). Dissociation between categorization and similarity judgment: Differential effect of causal status on feature weights. In U. Hahn & M. Ramscar, (Eds.), *Similarity and Categorization* (p. 87-107). New York, NY: Oxford University Press.
- Ahn, W.K., Kim, N.S., Lassaline, M.E., & Dennis, M.J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*, 361-416.
- Barr, R.A., & Caplan, L.J. (1987). Category representations and their implications for category structure. *Memory & Cognition*, *15*, 397-418.
- Barsalou, L.W. (1987). The Instability of Graded Structure: Implications for the Nature of Concepts. In U. Neisser (Ed.) *Concepts and conceptual development: Ecological and intellectual factors in categorization*. Cambridge: Cambridge University Press.
- Battig, W.F., & Montague, W.E. (1969). Category norms for verbal items in 56 categories: A replication and extension of the Connecticut Category Norms. *Journal of Experimental Psychology Monograph*, *80*, (3, Pt.2).
- Bloom, P. (1996). Intention, history and artifact concepts. *Cognition*, *60*, 1-29.
- Braisby, N.R. (1993). Stable concepts and context-sensitive classification. *Irish Journal of Psychology*, *14*, 426-441.
- Braisby, N.R. & Franks, B. (1997) 'What does word use tell us about conceptual content?', *Psychology of Language and Communication*, vol.1, no.2, pp.5-16.
- Braisby, N.R. & Franks, B. (2000). Perspectives and Categorization. Unpublished MS. Open University, UK.
- Braisby, N.R., Franks, B. and Harris, J. (1997) Classification and concepts: fuzzy or perspectival? In D. Dubois, (ed.) *Catégorisation et cognition: De la perception au discours* (pp). Paris: Kimé.
- Estes, Z. (2003). Domain differences in the structure of artifactual and natural categories. *Memory & Cognition*, *31*, 199-214.
- Hampton, J.A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language*, *34*, 686-708.

- Hampton, J.A. (1998). Similarity-based categorisation and fuzziness of natural categories. Cognition, 65, 137-165.
- Hampton, J.A., & Gardiner, M.M. (1983). Measures of internal category structure: A correlational analysis of normative data. British Journal of Psychology, 74, 491-516.
- Kalish, C.W. (1995). Essentialism and graded membership in animal and artifact categories. Memory & Cognition, 23, 335-353.
- Keefe, R., & Smith, P. (1997). Theories of Vagueness. In P.Keefe and P.Smith (Eds.) Vagueness: a Reader, (pp. 1-57). Cambridge: MIT Press.
- McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? Memory & Cognition, 6, 462-472.
- Medin, D.L., Lynch, E.B., Coley, J.D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? Cognitive Psychology, 32, 49-96.
- Murphy, G.L. (2002). The big book of concepts. Cambridge: MIT Press.
- Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. Psychological Review, 92, 289-316.
- Osherson, D.N., & Smith, E.E. (1997). On typicality and vagueness. Cognition, 64, 189-206.
- Rey, G. (1983). Concepts and Stereotypes. Cognition, 15, 237-262.
- Rips, L.J. (1989). Similarity, typicality and categorization. In S.Vosniadou & A.Ortony (Eds.), Similarity and analogical reasoning (pp. 21-59).
- Rosch, E., & Mervis, C.B. (1975). Family resemblances: studies in the internal structure of categories. Cognitive Psychology, 7, 573-605.
- Roth, E.M., & Shoben, E.J. (1983). The effect of context on the structure of categories. Cognitive Psychology, 15, 346-378.
- Smith, E.E., & Sloman, S. (1994). Similarity- versus rule-based categorization. Memory & Cognition, 22, 377-386.
- Weatherburn, C.E. (1961). A First Course in Mathematical Statistics. Cambridge: Cambridge University Press.

Appendix A

Items for categorization used in Experiments 1 - 3.

VEGETABLE	FRUIT	FISH	INSECT
spinach	strawberry	trout	ant
celery	orange	salmon	wasp
lettuce	banana	sardine	mosquito
asparagus	watermelon	catfish	grasshopper
potato	pomegranate	goldfish	termite
artichoke	tomato	shark	silkworm
soybean	avocado	eel	caterpillar
watercress	pumpkin	sea horse	moth
parsley	olive	squid	dust mite
dandelion	coconut	shrimp	head lice
seaweed	cucumber	jellyfish	maggot
bamboo shoot	eggplant	lobster	scorpion
chili pepper	acorn	starfish	centipede
cloves	almond	clam	spider
garlic	walnut	crab	tarantula
sage	date	tadpole	snail
apple	pine cone	whale	earthworm
turnip	rhubarb	seal	leech
peanut	sugar beet	plankton	tapeworm
bread	carrot	alligator	lizard
pineapple	mushroom	oyster	bat
milk	ginger	sponge	hamster
rice	onion	gull	amoebae
cereal	mint	frog	bacterium

Appendix A (continued)

SPORT	SCIENCE	TOOL	FURNITURE
swimming	chemistry	axe	chair
tennis	mathematics	screwdriver	table
skiing	astronomy	rake	bed
surfing	medicine	hammer	lamp
jogging	meteorology	shovel	desk
croquet	psychology	sewing needle	rug
billiards	nutrition	funnel	television
ballroom dancing	geometry	scalpel	shelf
frisbee	sociology	pitchfork	bookends
wrestling	mineralogy	calculator	curtains
darts	economics	dictionary	waste basket
hunting	geography	tractor	dishwasher
bullfighting	dentistry	toothbrush	cushion
weightlifting	pharmacy	broom	door mat
aerobics	architecture	scissors	painting
fishing	archaeology	key	ashtray
mountaineering	agriculture	varnish	telephone
hiking	criminology	screw	refrigerator
bridge	astrology	string	piano
kite flying	literature	umbrella	suitcase
conversation	advertising	photograph	plate
chess	palm reading	trunk	bucket
crosswords	religious studies	pen	pillow
picnicking	philosophy	stone	book

Appendix B

Examples of Pragmatic and Technical Context Scenarios used in all Experiments.

Technical Scenario for Vegetables.

The National Administration of Food and Agriculture is planning to regulate the growth of various kinds of agricultural produce, so that the quality of the produce in the market can be monitored.

Imagine that you belong to a panel of advisors for the Administration of Food and Agriculture to provide help in drafting the regulation. In the chapter for **vegetables**, you want to include all produce that should be considered as vegetables, excluding other kinds of agricultural produce, which would be covered under other chapters. Because the regulation affects vegetable farmers nationwide, a clear categorization of vegetables will thus ensure a fair and reasonable regulation. Consider each of the following items, and decide whether acting in the panel of advisors, you would classify the item in the category of **vegetables**.

Pragmatic Scenario for Furniture

Klein, which is a department store, is designing a sorting system to list the items in stock, as well as their prices and quantities. Marketing persons in the store can quickly obtain information about these items by using such a system. Imagine that you work for the department store to develop the sorting system. You have to categorize selling items under different headings, so that marketing persons can search for information easily and quickly. The following is a list of items from the department store. You have to decide whether or not they should be included in the category of **furniture**, so that most marketing persons would be able to find things under the category heading where they expected to find them.

Table 1

Percentage of Non-Modal Responses and Percentage of Categorization Responses Remaining the Same at Retest as a Function of Condition and Category for Experiment 1.

Category	Percentage Non-Modal Responses (NMR)				Percentage Categorization Responses Remaining the Same			
	No-Context	Pragmatic	Technical	Overall	No-Context	Pragmatic	Technical	Overall
Vegetable	17	15	14	16	92	90	92	92
Fruit	15	17	14	15	94	92	88	91
Fish	9	11	11	10	97	94	95	95
Insect	18	17	18	18	89	94	89	91
Sport	23	19	24	22	86	88	89	87
Science	26	23	24	24	88	90	85	87
Tool	24	17	20	20	87	90	86	87
Furniture	17	18	19	18	89	93	85	88
Overall	18.6	17.1	18.0	17.9	90	91	89	90

Table 2

Correlation of Categorization Probability with Typicality and Percentage of Positive Categorizations, as a Function of Condition and Category in Experiment 1.

Category	Correlation of Categorization with Typicality				Percentage of Positive Categorizations			
	No-Context	Pragmatic	Technical	Overall	No-Context	Pragmatic	Technical	Overall
Vegetable	.96	.97	.94	.96	53	53	57	53
Fruit	.94	.96	.97	.96	40	45	42	42
Fish	.95	.96	.95	.95	29	33	33	33
Insect	.98	.97	.99	.98	55	59	52	54
Sport	.87	.96	.91	.91	56	57	56	57
Science	.92	.92	.96	.93	54	61	60	61
Tool	.97	.86	.94	.94	52	45	50	51
Furniture	.95	.90	.96	.94	39	39	39	39
Overall	.94	.94	.95	.95	47	49	49	49

Table 3

Percentage of Non-Modal Responses (NMR), Correlation of Categorization Probability with Typicality, and Mean Categorization Probability as a Function of Condition and Category for Experiment 2, and for the No-Context condition of Experiment 1.

Category	Percentage of NMR		Correlation with Typicality		Categorization Probability		
	Pragmatic	Technical	Pragmatic	Technical	Pragmatic	Technical	No-Context (from Expt 1)
Vegetable	15	20	.95	.96	.51	.51	.53
Fruit	13	16	.95	.97	.38	.41	.40
Fish	16	19	.96	.97	.39	.41	.29
Insect	16	21	.97	.98	.46	.50	.55
Sport	20	26	.93	.96	.64	.51	.56
Science	19	26	.93	.93	.64	.60	.54
Tool	20	27	.89	.95	.47	.58	.52
Furniture	18	16	.91	.95	.42	.37	.39
Mean	17	21	.94	.96	.49	.49	.47

Table 4

Percentage of Extreme Responses, and Number of Participants (out of 20) Giving All Extreme Responses as a Function of Condition and Category for Experiment 3

Category	CONTEXT CONDITION			
	Pragmatic		Technical	
	Percent	Number	Percent	Number
Vegetable	84	5	84	7
Fruit	89	7	90	10
Fish	89	13	79	6
Insect	93	10	87	10
Sport	79	6	70	5
Science	86	9	82	7
Tool	82	8	78	3
Furniture	76	7	80	7
Mean	85	8.1	81	6.9

Figure Captions

Figure 1.

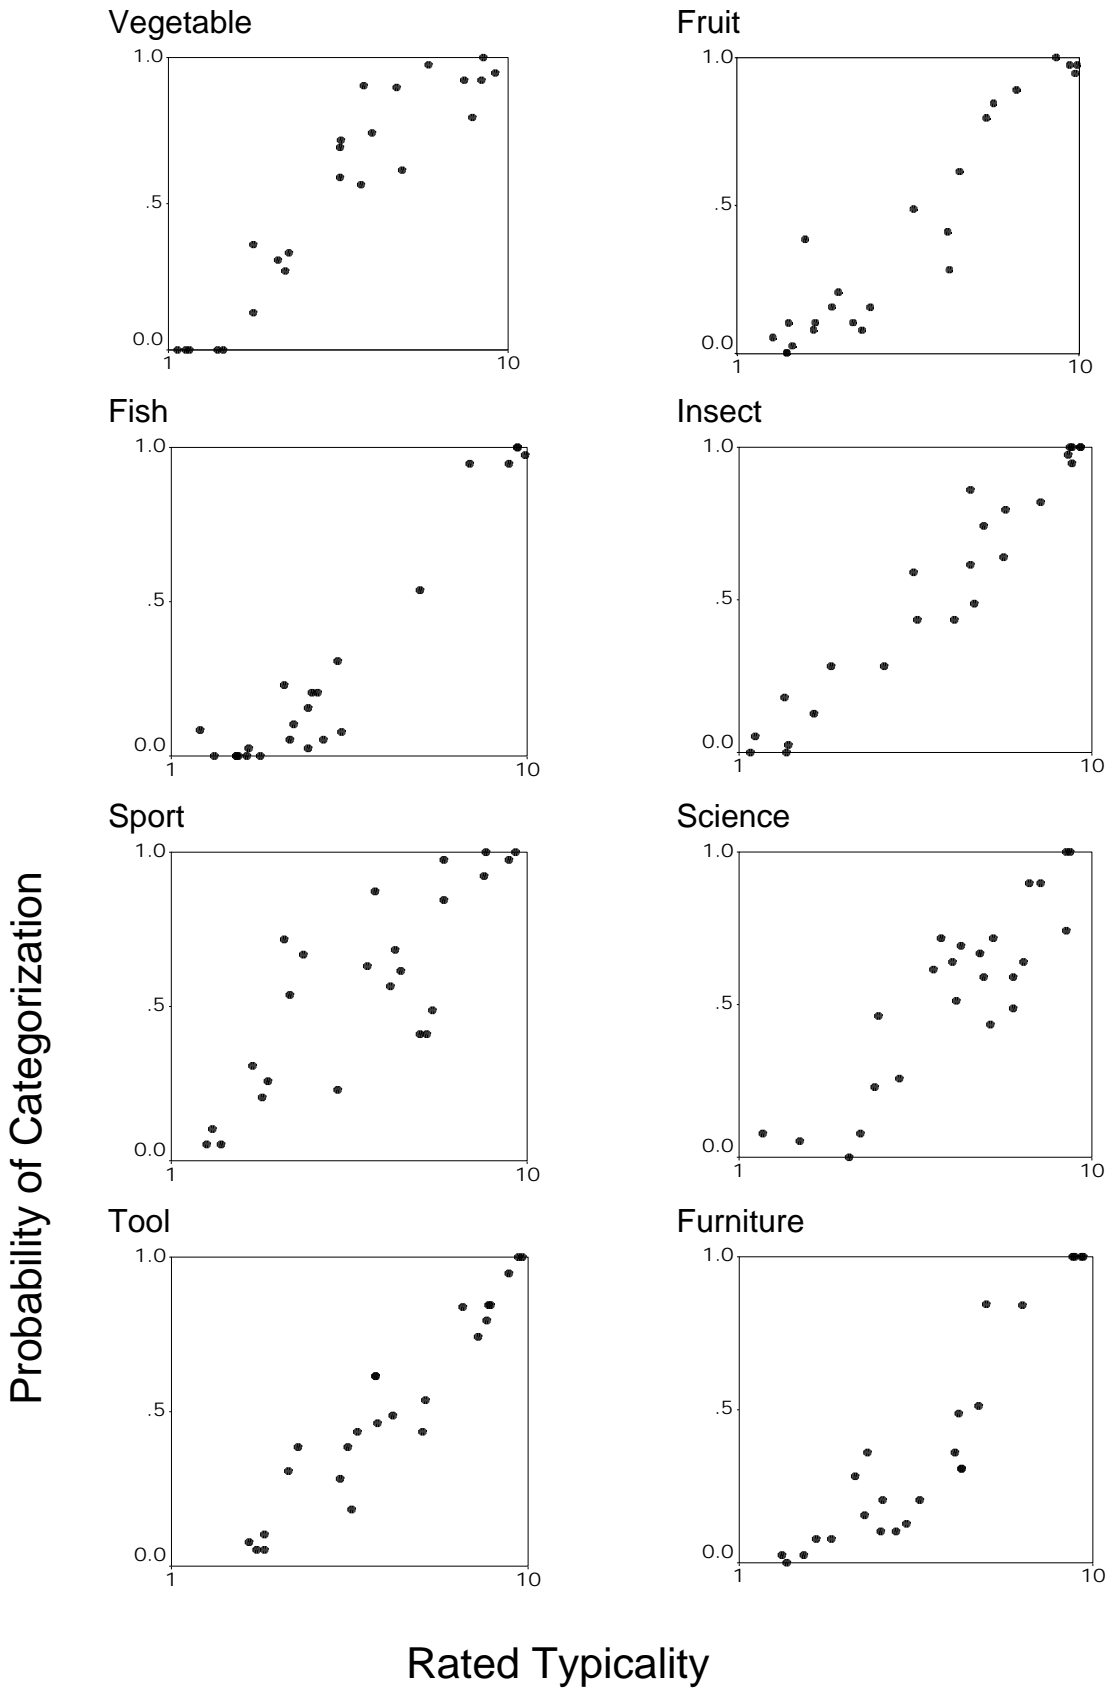
Scatter Plots of Categorization Probability with Typicality for Each Category in the No-Context Condition of Experiment 1.

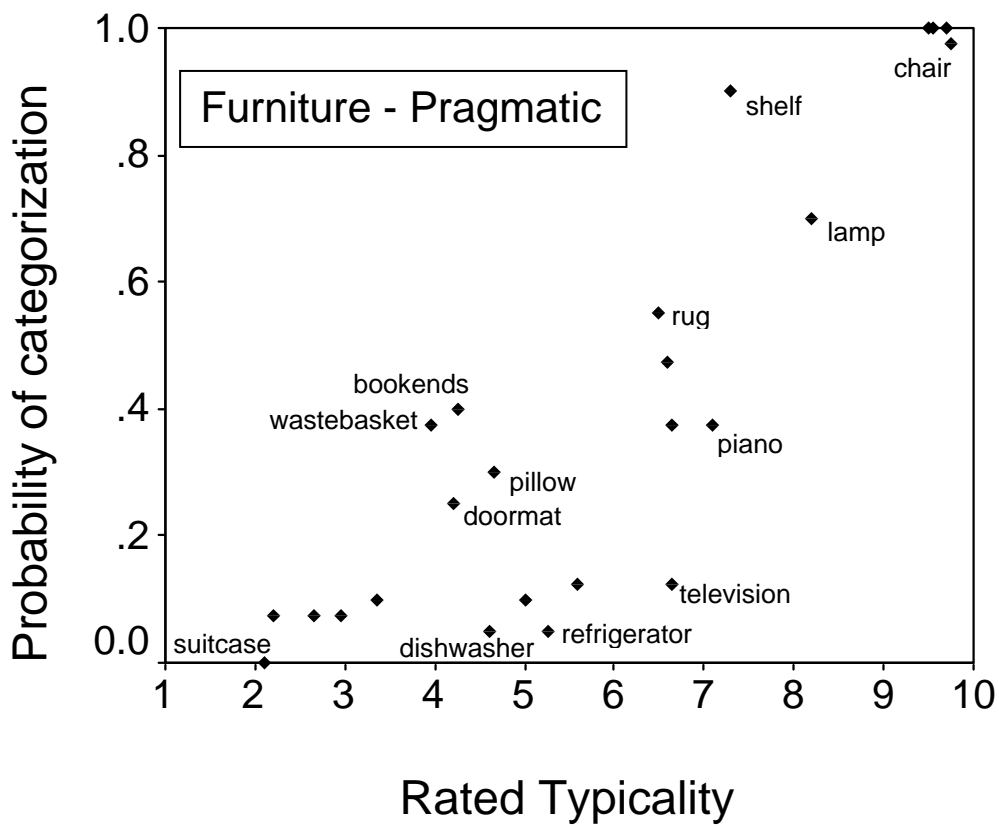
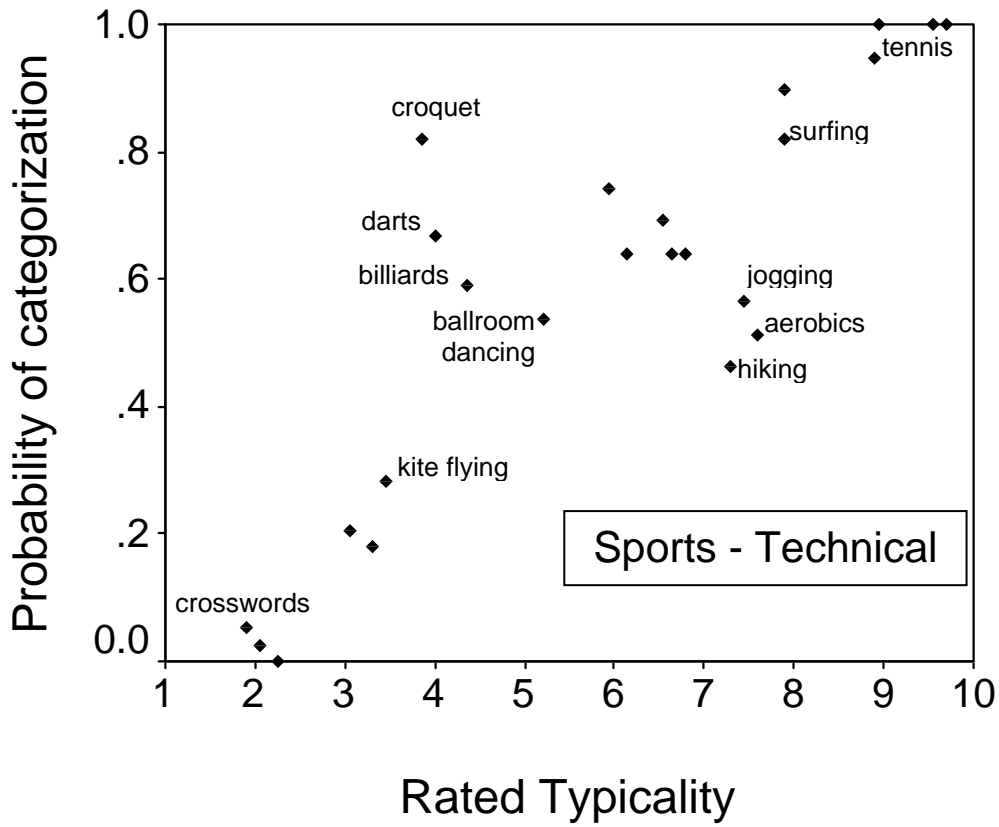
Figure 2.

Scatter Plots of Categorization Probability with Typicality for Sports in the Technical Context Condition, and for Furniture in the Pragmatic Context Condition in Experiment 1.

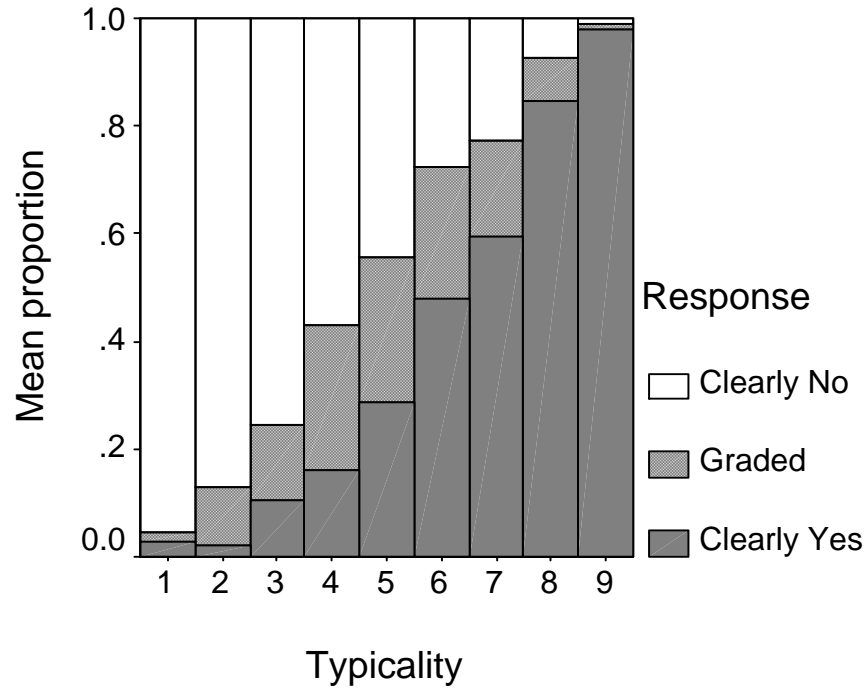
Figure 3

Stacked Bars Representing the Distribution of Responses for Each Condition in Experiment 3 Between Clearly Yes, Clearly No and Graded (1-7) Categorization Responses.





Pragmatic Condition



Technical Condition

