

9-1-2012

Locating the Source of Approach/Avoidance Effects on Natural Language Category Decisions

Matthew Zivot

University of Massachusetts - Amherst, mzivot@gmail.com

Follow this and additional works at: http://scholarworks.umass.edu/open_access_dissertations

Recommended Citation

Zivot, Matthew, "Locating the Source of Approach/Avoidance Effects on Natural Language Category Decisions" (2012). *Dissertations*. Paper 674.

This Open Access Dissertation is brought to you for free and open access by the Dissertations and Theses at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Dissertations by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

LOCATING THE SOURCE OF APPROACH/AVOIDANCE EFFECTS ON
NATURAL LANGUAGE CATEGORY DECISIONS

A Dissertation Presented

by

MATTHEW ZIVOT

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

September 2012

Psychology

LOCATING THE SOURCE OF APPROACH/AVOIDANCE EFFECTS ON
NATURAL LANGUAGE CATEGORY DECISIONS

A Dissertation Presented

by

MATTHEW ZIVOT

Approved as to style and content by:

Andrew Cohen, Chair

Charles Clifton, Jr., Member

Ronnie Janoff-Bulman, Member

Angela De Oliveira, Member

Melinda Novak, Department Head
Psychology

ACKNOWLEDGEMENTS

I would like to thank my advisor, Andrew Cohen, for his insights, suggestions and willingness to go to bat for me. Thanks also to Chuck Clifton, who welcomed me into the program and helped me to think through the field.

Thanks to Mindi Rock for selflessly sharing her data and to Lisa Fiorenzo for selflessly helping me collect my data.

I am forever indebted to my wife, Emily Zivot. From helping organize these thoughts through helping me edit this manuscript, she supported me through this entire dissertation. Thank you for learning to distinguish an exemplar model from a prototype model just so you could help me.

Finally, thanks and love to my parents. For years of unconditional support, all they wanted in return was to attend my defense.

ABSTRACT

LOCATING THE SOURCE OF APPROACH/AVOIDANCE EFFECTS ON NATURAL LANGUAGE CATEGORY DECISIONS

SEPTEMBER 2012

MATTHEW ZIVOT, B.A., THE GEORGE WASHINGTON UNIVERSITY

M.S., UNIVERSITY OF MASSACHUSETTS AMHERST

PH.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Andrew Cohen

In this dissertation, two exemplar-based models of categorization, the General Context Model (GCM) and the Exemplar Based Random Walk model (EBRW), were used to describe between-group categorization differences in artificial and natural language categories. Prior research has shown that political Conservatives in avoidance mode are more exclusive categorizers of natural language category members than Conservatives in approach mode, but this effect was absent for Liberals (Rock & Janoff-Bulman, 2010). In Experiment 1, experimenter-generated stimuli were used to show that the EBRW could account for between-group differences in categorization decisions. In Experiment 2, the data collected by Rock and Janoff-Bulman were used to develop techniques allowing the GCM to account for between-group differences in natural language categorization decisions. Experiment 3 extends these methods to allow the EBRW to account for between-group differences in natural language categorization decisions. Across these experiments, the models identify between-group differences in determining similarity, bias to give an “in-the-category” decision, and the amount of information required to make a categorization decision. Techniques for modeling natural language categorization decisions are discussed

CONTENTS

	Page
ACKNOWLEDGEMENTS.....	iv
ABSTRACT.....	v
LIST OF TABLES	viii
LIST OF FIGURES	ix
CHAPTER	
1. INTRODUCTION	1
2. LITERATURE REVIEW	5
The Process of Categorization	5
Modeling the Categorization Process	15
Modeling Natural Language Categorization.....	26
Locating the Effects of Mode in the EBRW	31
3. OVERVIEW OF EXPERIMENTS	35
4. EXPERIMENT 1	37
Method	38
Procedure	41
Results.....	43
Discussion.....	48
5. EXPERIMENT 2	52
Method	54
Results.....	61
Additional Modeling Considerations.....	65
Discussion	69
6. EXPERIMENT 3	73
Method	73
Results.....	75
Discussion.....	83

7. GENERAL DISCUSSION	99
The Effect of Mode and Political Identity	99
Modeling Natural Language Categories	103
Conclusion	105
APPENDICES	
A. PRE-SCREENING SURVEY QUESTIONS.....	106
B. INSTRUCTIONS FOR APPROACH/AVOIDANCE MANIPULATION IN EXPERIMENTS 1 AND 3.....	107
C. STIMULI FOR EXPERIMENT 1	108
D. DIMENSIONAL ANALYSIS FOR MULTIDIMENSIONAL SCALING.....	109
E. WORD ASSOCIATIONS FOR NATURAL LANGUAGE CATEGORY MODELING	111
F. RESULTS FROM EXPERIMENTS.....	115
REFERENCES	132

LIST OF TABLES

Table	Page
1. Average %yes Judgments (and SE) by Manipulation and Stimulus Type	44
2. Average Parameter Values (and SE) for Signal Detection Analysis	46
3. Optimal Number of MDS Dimensions and Correlation with Measured Typicality by Category	60
4. Results from Best Fitting Prototype Models	68
5. Average “In-the-category” Rates for Experiment 3	77
6. Logistic Regression for Experiment 3	78

LIST OF FIGURES

Figure	Page
1. Example of GRT category boundaries	18
2. Example of the random walk process in the EBRW	24
3. Example of a categorization decision as described by the EBRW	32
4. Example jellyfish stimuli	40
5. Average and model predicted percent “in-category” judgments by stimulus	49
6. Average and model predicted RT by stimulus	49
7. Figure 1 from Rock and Janoff-Bulman	55
8. Best fitting model predictions for Clothing alongside participants’ responses	62
9. Best fitting model predictions for Furniture alongside participants’ responses	63
10. Best fitting model predictions for Vehicles alongside participants’ responses	64
11. Experiment 3 categorization rates by typicality as predicted by the logistic regression	79
12. Furniture “in-the-category” judgments and modeling predictions for Conservatives in avoidance mode and all other participants.	84
13. Furniture RT and modeling predictions for Conservatives in avoidance mode and all other participants.	87
14. Vehicle “in-the-category” judgments and modeling predictions for Liberals and Conservatives	90
15. Vehicle RT and modeling predictions for Liberals and Conservatives	93

CHAPTER 1

INTRODUCTION

Sometimes, the Supreme Court must step in to make a categorization decision for a nation torn apart. In 1893 it officially decided that tomatoes are Vegetables and not Fruit in *Nix v. Hedden* (1893). In the Court’s decision the common use of tomatoes—as an ingredient in salads and main dishes rather than desserts—is taken as evidence that tomatoes are Vegetables. This decision, however, has not ended the debate. Just recently Ohio named the tomato its State Fruit. The Supreme Court may have to revisit this issue in the near future.

Why do these debates continue? What is different about people who categorize tomatoes as fruits and those who categorize them as vegetables? Researchers have identified many causes of between-group differences in categorization decisions, including goals (e.g., Barsalou, 1983), mood (e.g., Isen & Daubman, 1984), background knowledge (e.g., Lin & Murphy, 1997), and the interaction of political identity and being in approach or avoidance mode (Rock & Janoff-Bulman, 2010). Researchers are still uncertain, however, which step or steps in the categorization process are influenced by these contextual factors.

This dissertation will use models of categorization to account for between-group differences in categorization observed by Rock and Janoff-Bulman. They found that when Conservative participants were placed into avoidance mode—a state in which people focus on preventing negative outcomes—they became more exclusive categorizers than when Conservative participants were placed into approach mode—a state in which people focus on bringing about positive outcomes. That is, Conservatives in avoidance mode required items to be highly typical of their category before they made an “In-the-category” judgment. Liberals, however, were found to be equally inclusive regardless of whether they were in approach or avoidance mode.

In order to locate the source of between-group differences in categorization decisions, one must first identify the stages of categorization. For millennia, categorization was considered to be a rule-defined process in which people determined whether a to-be-categorized item met all the criteria set forth by the rules that govern a category (Murphy, 2002). For example, a tomato is a fruit if it has seeds within a fleshy covering. Theoretical work by Wittgenstein (1953) and empirical studies by Rosch and Mervis (1975) have provided evidence that categories are not rule-based but instead are family-resemblance based. For example, a tomato has uses in common with many vegetables and therefore should be grouped with the items to which it is most similar. According to this view, the process of categorization depends on determining the dimensions along which category members are similar and how this similarity is translated into category membership (Medin & Schaffer, 1978).

Many theories of categorization focus on how the features of a to-be-categorized item impact a categorization decision (for a critique, see Murphy & Medin, 1985). As discussed, though, factors other than the item itself have been shown to influence category related decisions. For the purposes of this dissertation, these factors will be referred to as context.

Researchers who have investigated components of the categorization process that are affected by context have focused on the selection of pertinent dimensions (e.g., Kelemen & Bloom, 1994; Lin & Murphy, 1997). (For example, attending to the color of a person's clothing may be essential for categorizing friends versus foes during a pick-up basketball game but bears little weight while ice skating.) Researchers have specifically designed their experiments to identify differences in this component of categorization. While determining how attention is distributed across dimensions is an important component of the categorization process, it is neither the only component of the categorization process nor the only one in which context has an effect. For instance, researchers have found that approach and avoidance mode not only impacts how attention

is distributed across dimensions (Förster, Friedman, Özelsel, & Denzler, 2006; P. A. Gable & Harmon-Jones, 2008), but also whether commonalities or differences are given preference for making similarity judgments (Förster, 2009) and how decision boundaries are established (Markman, Baldwin, & Maddox, 2005). If context can affect components of categorization other than the selection of pertinent dimensions, more sensitive experimental methods will be required to capture these effects.

Process models of categorization like SUSTAIN (Love, Medin, & Gureckis, 2004) and the General Context Model (Nosofsky, 1984; Nosofsky, 1986) represent multiple components of the categorization process. Such models can be used to identify any component in which context creates between-group categorization differences so long as the model has a parameter to describe that component. In order to describe between-group differences, a model must be selected that has parameters for components of categorization that are affected by context.

Exemplar models, such as the General Context Model and the Exemplar Based Random Walk (Nosofsky & Palmeri, 1997), are good candidates for identifying the components of the categorization process that are affected by context. An exemplar-based model represents a category with a set of previously seen items whose category membership is known (i.e., exemplars). In these models a to-be-categorized item is compared to exemplars and categorized based on the category membership of the exemplars to which it is most similar. A prototype-based model represents a category with an abstraction of previously seen items whose category membership is known (Hampton, 1993). Exemplar-based models may be better able to describe the effects of ephemeral contexts than prototype-based or rule-based models because information about previously seen category members is preserved. Furthermore, exemplar models have parameters that can be matched with cognitive processes upon which approach and avoidance modes have an impact.

The purpose of this dissertation, therefore, is to determine whether the General Context Model and the Exemplar Based Random Walk can identify the source of between-group differences in categorization decisions reported by Rock and Janoff-Bulman (2010). To do so requires:

1. Using the EBRW to examine between-group differences in categorizing artificial stimuli;
2. Extending the GCM to account for between-group differences in natural language categorization; and finally,
3. Extending the EBRW to account for between-group differences in natural language categorization.

CHAPTER 2

LITERATURE REVIEW

The Process of Categorization

If a category can be defined with a set of necessary and sufficient rules, then categorization is a simple task. When an item fulfils those rules, it is a member of the category; when it does not, it is not. Socrates is looking for just such a set of rules on the way to his trial when he says, “I did not bid you to tell me one or two of the many pious actions, but the form itself that makes all pious actions pious” (Plato, 1975). Philosophers spent millennia trying to define necessary and sufficient rules for categories, but could only do so for the most artificial of categories like the mathematical category: prime numbers.

In the middle of the twentieth century, Wittgenstein (1953) concluded that categories are not based on necessary and sufficient rules, but are based on family resemblance. He writes that family resemblances are “a complicated network of similarities overlapping and crisscrossing.” For example, as a Dead Head you might listen to jam music, wear tie-dye, go on tour with a band, and sell plates of rice and beans in a parking lot, but you may not have to do all of these things to be a Dead Head.

If categories are based on family resemblance, membership can be a continuous measure rather than an all-or-nothing state. Evidence for this continuum has come from different measurements. Rosch (1975) showed that different people tended to agree on how typical an item was of a given category. Rips, Shoben, and Smith (1973) showed that an item’s typicality as a member of a given category predicts the speed with which a person can make a categorization decision about that item. McCloskey and Glucksberg (1978) showed that an item’s typicality as a member of a given category predicts the & Blok, 2005) or average (Hampton, 1993) representations of categories. For example, the prototypical representation of Bird might have all of the features associated with birds, such as flying, having feathers, and living in trees. It also would lack all of the features

that are associated with other categories, such as swimming, having scales, or nesting on the ground. According to proponents of exemplars, a to-be-categorized item is compared to a set of previously seen items that represent a category (Medin & Schaffer, 1978). For example, an item seen moving across the sky might be compared to specific birds, planes, and super-heroes that you have previously seen.

The Effect of Context on Categorization

Although category membership is on a continuum, there is no reason for category decisions to be inconsistent over time (McCloskey & Glucksberg, 1978) unless the context within which a category decision is made matters. An ephemeral context like a temporary goal or stance may account for some of the intrapersonal variability in categorization decisions observed by McCloskey and Glucksberg. Barsalou (1983) showed that participants provided with a goal that two items—normally considered dissimilar—can both accomplish were more likely to rate those items as similar. For example, when participants were given the goal of getting a birthday gift, they rated CDs and necklaces as more similar than participants who were not given a goal. Lombrozo (2009) showed that participants primed to take a teleological stance (focusing on ends) were more likely to say that items were members of a given category if those items had the same parts as known category members. Participants primed to take a mechanistic stance (focusing on means) were more likely to say that items were members of a given category if those items behaved in the same way as known category members. For example, a teleological stance might make you more likely to accept a person wearing a suit who listens to the Grateful Dead as a Dead Head, but a mechanistic stance might favor you towards a person wearing tie-die who does not listen to the Grateful Dead.

An ephemeral context like a temporary mood or mode may account for intrapersonal variability in categorization decisions. Ell, Cosley, and McCoy (2011) showed that participants under stress were able to identify members of a learned category more accurately than participants who were not under stress. Isen and Daubman (1984)

showed that participants who were given a free gift or shown a funny movie were more willing to include atypical category members in a given category (e.g., olives in Fruit) than those who did not receive a free gift or who saw a boring movie. Using a similar categorization task, Price and Harmon-Jones (2010) showed that mode (approach vs. avoidance) may alter the inclusiveness of categorizers who are in a good mood. Participants were asked to smile, emulating a positive mood, and then to lean forward, sit upright, or recline, representing a range from high to low approach mode. Participants in the low-approach mode position (i.e., recliners) were more inclusive when categorizing atypical items than participants in the high-approach mode position (i.e., leaners).

Context like accessories that are incidental to a to-be-categorized item—ephemera that could change while the to-be-categorized item itself does not—may account for some of the intrapersonal variability in categorization decisions. Macrae, Bodenhausen, and Milne (1995) showed that participants shown a video of an Asian woman eating with chopsticks were more likely to categorize her as Asian while participants shown video of the same woman applying make-up were more likely to categorize her as a Woman.

Finally, categories can be imagined as fitting into a taxonomy, with more abstract superordinate labels and less abstract subordinate labels. For instance, a sparrow is a Bird, which is an Animal. Ephemeral contexts may account for some of the intrapersonal variability in categorization decisions by changing the way that people use this taxonomy when categorizing. Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976) and Collins and Quillian (1969) showed that categories often display what is termed a “basic level effect.” Participants in both studies displayed a preference for a certain level of abstractness when categorizing. For example, an item was categorized as a Chair first and faster than as Furniture (superordinate) or as an Easy Chair (subordinate). Ephemeral context like the scene that surrounds a to-be-categorized item can diminish or erase the basic level effect (Murphy & Wisniewski, 1989). For example, an object might be categorized as a Settee when it is selected from a showroom but may be categorized as a

Couch when it is offered as a seat to a friend. In part, this is due to the other objects from which the to-be-categorized item needs to be distinguished.

Rosch (1975) showed that when participants were primed with a superordinate category label, they read words representing typical members of that category faster and words representing atypical members of that category slower than when they were not primed with a superordinate category label. Temporarily changing the context in which a superordinate category label is seen changes the facility with which participants read words representing typical and atypical members of that category (Roth & Shoben, 1983). For example, the word “goat” was read faster than the word “horse” following a sentence about “milking the animal” even though horses are more typical members of the category Animal than goats.

Some contexts, like background knowledge, are more persistent and may account for interpersonal differences in categorization decisions. Tanaka and Taylor (1991) showed that the extent of background knowledge participants had about an item changed the basic taxonomic level at which that item was categorized. Experts identified items at a more subordinate basic level than novices. For example, a bird expert who spots a scarlet tanager is likely to categorize it as a Tanager, while a non-birder is likely to categorize it as a Bird¹. Yamauchi and Yu (2008) showed that adding background knowledge changed the types of inferences participants made about the members of a category. When participants were told that an item’s label reflected category membership they inferred that the item shared features of prototypical category members. When they were told that an item’s label reflected the place where it lived or the type of food it ate they inferred that the item shared features of the most similar exemplar that they had seen.

Persistent contexts like group identity may also affect categorization decisions. Medin, Cox, and Ross (2006) showed that sport and subsistence fishermen in the same region rate the typicality of local fishes differently even though both groups were highly

¹ Note: This is taxonomically equivalent to categorizing squirrels, dogs, and humans as Mammals, so perhaps—not surprisingly—we are all mammal experts.

familiar with these fish. Lynch, Coley, and Medin (2000) showed that tree taxonomists and horticulturalists both used goals for determining tree typicality, but differences in their goals led them to rate the typicality of trees differently. Ratneshwar, Barsalou, Pechmann, and Moore (2001) showed that participants identified as healthy eaters and participants who were not identified as healthy eaters rated the similarity of candy bars and granola bars differently when they were asked to consider foods one might eat while driving. Lombrozo (2009) showed that participants who naturally take a teleological stance toward the world were likely to include items with similar purposes in the same category, while participants who naturally take a mechanistic stance toward the world were more likely to include items with similar constructions in the same category. Rock and Janoff-Bulman (2010) showed that participants who self-identified as politically conservative were more likely to exclude items of moderate typicality from a given category when those participants were put in avoidance mode. For example, they were less likely to say that a bookcase is Furniture or that a yacht is a Vehicle. Participants who self-identified as political liberals were not more likely to exclude items of moderate typicality from a given category when those participants were put in avoidance mode.

Finally, persistent context like the relationships between category members can affect categorization decisions. Categories can be more or less coherent based on how well they “go together in light of prior theoretical, causal, and teleological knowledge” (Patalano, Chin-Parker, & Ross, 2006). Patalano, Wengrovitz, and Sharpes (2009) and Patalano, Chin-Parker, and Ross (2006) showed that people are more likely to make inferences based on an item’s membership in a more coherent category than that item’s membership in a less coherent category. Participants were more likely to make inferences about feminist waiters based on their membership in the category Feminists than in their membership in the category Waiters.

Identifying the Effect of Context on Attention to Dimensions

Experimenters that attempt to identify specific components of the categorization process that are affected by context often use artificial stimuli or artificial categories. These stimuli have high internal validity (but low ecological validity) because experimenters create them to operationalize the specific theories about categorization that they are investigating. For example, Yamauchi and Yu (2008) created “bug” stimuli with five different dimensions (antennae shape, head shape, number of legs, thorax markings, tail shape), each of which had two feature values. The category label assigned to each variation of these stimuli was determined based on family resemblance theories of category structure. These clearly delineated dimensions with easily distinguishable features are designed allow researchers to identify how these dimensions are used for the purpose of categorization.

For example, Kelemen and Bloom (1994) designed stimuli that were circles of different colors and sizes. When participants were told that they were categorizing Tiny Machines they categorized primarily along the dimension of size; when participants were told that they were categorizing Microscopic Animals they categorized primarily along the dimension of color. Background knowledge about Animals and Machines is likely to have changed the dimensions that participants used when making categorization decisions.

A series of studies using the novel categories of Arctic Vehicles and Jungle Vehicles have shown that background knowledge affects the selection of pertinent dimensions during categorization decisions (Kaplan & Murphy, 1999; Kaplan & Murphy, 2000; Murphy & Allopenna, 1994; Spalding & Murphy, 1999). When making decisions about novel categories, background knowledge made it easier for participants to build multi-dimensional family resemblance categories rather than one-dimensional rule-based categories (Kaplan & Murphy, 1999). It also made it easier for participants to learn categories that require attention across multiple dimensions (Kaplan & Murphy,

2000; Murphy & Allopenna, 1994) without reducing their sensitivity to feature frequency (Spalding & Murphy, 1999). In all of these studies, participants used more of the dimensions of the vehicles for their categorization decisions than when the dimensions did not allow them to apply their background knowledge.

Sometimes researchers create both artificial stimuli and artificial categories to pinpoint a step in the categorization process that they think context affects. For example, Lin and Murphy (1997) created an artificial category Tuk and artificial stimuli with four dimensions (a loop, a cone, a tube, and a string), each of which could have a value of present or absent. One group of participants learned that a Tuk was an animal catching device; a second group of participants learned that a Tuk was a pesticide spreading device. Participants with different background knowledge paid attention to different dimensions when making categorization decisions about which artificial stimuli were in the artificial category Tuk and which were not.

Effects of Approach/Avoidance Mode on Cognition

The purpose of this dissertation is to account for the between-group differences in categorization decisions identified by Rock and Janoff-Bulman (2010). Namely, that Conservatives were more likely to exclude items of moderate typicality from a given category when they were put in avoidance mode. Research on approach and avoidance modes and cognition indicate several potential sources of this effect: spreading attention (Friedman & Förster, 2005; Förster et al., 2006; P. A. Gable & Harmon-Jones, 2008; P. Gable & Harmon-Jones, 2010; Koch, Holland, Hengstler, & van Knippenberg, 2009), attention to similarities as opposed to dissimilarities (Förster, 2009), and changing decision criteria (Markman et al., 2005). In each of the studies discussed in the previous section, researchers have investigated the effect of context on one component of the categorization process: the selection of pertinent dimensions. The current section will show that accounting for the between-group differences in categorization by people

in approach and avoidance mode will require a method sensitive to differences in components other than the selection of pertinent dimensions.

Approach and avoidance mode may impact the spread of attention though the direction of its effect is unclear. Participants have a narrow spread of attention when they focus on component elements (i.e., the trees); they have a broad spread of attention when they focus on composite elements (i.e., the forest). Förster et al. (2006) primed participants for approach or avoidance and then had them respond to composite letters that were comprised of smaller component letters (called Navon letters). They found that participants in the approach condition were faster to respond to the composite letters and participants in the avoidance condition were faster to respond to the component letters. This indicates that participants in the approach condition were spreading their attention more broadly to see the larger picture and participants in the avoidance condition were narrowing their attention to see the details.

Approach and avoidance mode may interact with emotion to affect the spread of attention. In a study of mode and positive affect (P. A. Gable & Harmon-Jones, 2008), participants were primed for either low-approach mode or high approach mode and then asked to choose the best pair from a group that matched either on composite shape or component parts. Gable and Harmon-Jones found that participants in the low approach condition selected shapes with composite similarity more often than participants in the high-approach condition. In a similar study, Gable and Harmon-Jones (2010) primed participants with low levels of avoidance or high levels of avoidance and asked them to respond to Navon letters. They found that participants in the low avoidance condition were faster to identify composite letters than component letters, while participants in the high avoidance condition were faster to identify component letters than composite letters. In contrast to the findings of Förster et al. (2006), these studies indicate that both high avoidance and high approach modes lead to a narrower distribution of attention relative to low avoidance or low approach modes.

The Stroop task is another measure that approach and avoidance modes impact, and again the direction of that impact is unclear. Friedman and Förster (2005) found that participants primed with approach mode were faster to respond to incompatible color trials than participants primed with avoidance mode. In a contrasting study using embodied cognition techniques, Koch, Holland, Hengstler and van Knippenberg (2009) found that when participants were primed with avoidance mode, they were faster on the incompatible color trials than when they were primed with approach mode. Not only do these results contradict each other, but the authors' interpretation of them may conflict as well. Koch et al. take their results as evidence that, "avoidance cues facilitate the recruitment of cognitive control." Friedman and Förster, however, take their results to mean that approach mode serves to, "enhance attentional flexibility."

In the end, what is being "cognitively controlled" is which cues to attend to and which to ignore, so all of these studies imply that approach and avoidance modes impact the distribution of attention. This explanation holds regardless of whether approach mode broadens attention, avoidance mode broadens attention, or they both decrease it. The effect still lies at the distribution of attention. Therefore, any method used to identify between-group differences in categorization performance due to approach and avoidance mode should have a way of measuring the distribution of attention.

Recent research has reinterpreted the meaning of the broad versus narrow spreading of attention as attending to similarities as opposed to dissimilarities. Förster (2009) found that, when asked to compare two items, participants who were primed with broad attention listed more similarities than dissimilarities while participants primed with narrow attention listed more dissimilarities than similarities. This implies that approach and avoidance modes change the way that people weight similarity, either focusing participants on what makes items more similar to each other or more dissimilar to each other. Therefore, any method used to identify the source of the differences between modes should be able to account for differences in similarity weighting.

Research on the interaction between approach and avoidance mode and category structure during a category learning task indicates that mode can influence a decision criterion Markman, Baldwin and Maddox (Markman et al., 2005) found that when participants primed with approach mode learned a one-dimensional category with an approach feedback structure, they used a categorization criterion that optimized reward rather than performance. Participants primed with approach mode learning a one-dimensional category with an avoidance feedback structure used a criterion to optimize performance rather than reward. Conversely, participants primed with avoidance mode used a criterion to optimize reward with an avoidance feedback structure but optimized performance with an approach feedback structure. In this one-dimensional category structure, spread of attention across dimensions could not be a factor. This finding implies that a method used to identify between-group differences due to approach and avoidance modes should be able to identify changes in decision boundary as well.

Finally, Rock and Janoff-Bulman (2010) found that conservatives in avoidance mode were less likely to include moderately typical items in a category, while Price and Harmon-Jones (2010) found that participants in avoidance were more likely to include highly atypical items in a category. This difference in results may be explained by changes in consistency of categorization decisions. If avoidance mode induces probability matching, categorizers in avoidance mode would make “in-the-category” decisions in proportion to their perceived probability that the to-be-categorized item is a category member, as opposed to consistently giving the most likely response. Practically speaking, this behavior would increase the inclusion of moderately atypical items while depressing the inclusion of moderately typical items. Any method used to identify the source of approach/avoidance differences on categorization should include a way to measure the consistency of responding. Research by Janoff-Bulman, Sheikh and Baldacci (2008) indicates that political conservatism tend toward avoidance mode while political

liberalism tend toward approach mode, so a method that can capture the effects of approach/avoidance should be able to capture the effects of political identity as well.

Modeling the Categorization Process

The previous section shows that approach and avoidance modes affect the spread of attention across dimensions, whether commonalities or dissimilarities are attended to and how decision criteria are established. It is clear that a method for locating the components in the categorization process affected by approach/avoidance will need to capture effects other than how attention is distributed across dimensions. A promising method for identifying the effects of approach/avoidance is to use models of categorization. Process models formalize the components of the categorization process, specifying parameters that describe each component (Kruschke, 2008). Differences in these parameters can reflect between-group differences in specific components of categorization.

When formalizing the categorization process, models make assumptions about whether categories are represented with prototypes or exemplars. They also break up the categorization process into different components. The assumptions that a model makes about the process of categorization affects its ability to capture the between-group differences caused by approach and avoidance mode. Some models of categorization, such as ATRIUM (Erickson & Kruschke, 1998) and ALCOVE (Kruschke, 1992) have been shown to perform well at capturing categorization decisions. Since they are implemented as connectionist models that slowly learn categories over time, however, they are not well suited to capturing the spontaneous changes in categorization behavior observed in Rock and Janoff-Bulman (2010). Other models, such as the Rational Model (Anderson, 1991) use a Bayesian analysis to describe optimal categorization performance. This would not be well suited to identifying the reason for between-group differences, since there can be only one optimal categorization strategy. This section will survey some of the prominent models of categorization that either explicitly address

between-group differences or have been designed in a way to capture the impacts of approach and avoidance mode.

Knowledge Resonance

The Knowledge Resonance Model is a connectionist prototype model that has the effect of background knowledge built in (Rehder & Murphy, 2003). It represents a to-be-categorized item's features as input nodes and category labels as output nodes. For example, a connectionist model of Dead Head categorization could have an input node for tie-die shirts. When this input node is activated in the presence of a tie-die shirt, it excites the output node of Dead Head. The Knowledge Resonance Model (KRES) is constructed to represent the kind of association between features that background knowledge brings to the categorization process (e.g., Hoffman, Harris, & Murphy, 2008; Murphy & Allopenna, 1994). For example, knowing that a person wears tie-die magnifies the effect of knowing that they tour with a band when exciting the Dead Head category label. Furthermore, it mutes the potential activation of contrasting features that could inhibit the activation of the Dead Head category label, such as living in a Manhattan penthouse. KRES represents these relations with excitatory connections between features that are associated with each other and inhibitory connections between features that associated with a contrasting category.

KRES is unsatisfactory for capturing the effects of approach and avoidance on categorization because the effects of context are built into the model. KRES represents between-group differences as the excitatory and inhibitory associations between features. These relations are not flexible and cannot reflect the ephemeral changes that occur when a person changes goals, moods or modes. Furthermore, KRES does not break down categorization in a way that provides analogs with the aspects of cognition that approach and avoidance mode have been shown to impact, such as attending to similarities versus dissimilarities and changing decision boundaries. All effects are all conflated in the weights between nodes.

Baywatch

A second model that captures the effect of background knowledge is called Baywatch (Heit, Briggs, & Bott, 2004). It is also a connectionist prototype model in which input feature nodes get associated with output category labels. The model is modified so that features also can get associated with latent concepts, which themselves get associated with category labels. For instance, if you did not know the concept of Parrot Heads, you would be unlikely to apply that category label to people wearing flip-flops, drinking margaritas and following Jimmy Buffet. Once you realize that Parrot Heads are similar in concept to Dead Heads, wearing flip-flops and drinking margaritas gets associated with concept of Dead Head. The latent concept of Dead Head then gets associated with the category label Parrot Head, providing extra activation between Parrot Head input nodes and the Parrot Head category label.

While Baywatch may partially represent the impact that background knowledge has on categorization decisions, it cannot be extended to account for the effects of ephemeral context observed by Rock and Janoff-Bulman (2010). Like KRES, the effect of background knowledge is built into the model. In Baywatch this effect is included as the presence of latent concepts, preventing this model from capturing ephemeral context effects on categorization. Also like KRES, all effects in Baywatch are captured in the weights between input nodes, latent concepts, and output nodes. This makes Baywatch unable to identify components of the categorization process that are affected by approach/avoidance mode.

General Recognition Theory

General Recognition Theory is a modern version of rule-based categorization theories (Ashby & Maddox, 1993; Maddox & Ashby, 1993). For instance, imagine a group of Dead Heads arguing about how to determine category membership. One is only willing to include people who have seen a minimum number of shows. Another is only willing to include people who have accumulated a minimum number of taped

concerts. A third allows a more generous combination of the two – if you have seen some shows and possess some taped concerts, you are categorized as a Dead Head. The General Recognition Theory (GRT) represents the categorization process by fitting a decision boundary that separates members of different categories.

For the first categorizer, it will locate a linear boundary along the “shows-attended” dimension, for the second categorizer it will locate a linear boundary along the “taped-concerts” dimension, and for the third categorizer it will find a combination of both dimensions that best accounts for their categorization decisions. Figure 1 shows a graphical representation of these boundaries.




Figure 1: Example of GRT category boundaries

The GRT does a good job of capturing the spread of attention across dimensions. It could capture the effects of approach and avoidance observed by Forster et al. (2006) and Gable and Harmon-Jones (2008). Narrowed attention could be represented by a one-dimensional rule while broad attention could be represented by a combination across multiple dimensions. The GRT, however, would not be able to capture potential differences in whether commonalities or dissimilarities are given preference since it has no way of describing dimensions in this fashion. The GRT also has no way to determine if categorization decisions have become more or less consistent, since there are no regions of uncertainty around the decision boundary. It predicts that every item on a given side of the boundary will be included or excluded with equal probability.

The GRT has been used to identify between-group differences in categorization. Ell et al. (2011) had participants learn to categorize a two-dimensional stimulus in a way that required attending to both dimensions after they had been placed in a stressful situation. The researchers used the GRT to determine that a participant's perception of threat was associated with their using an appropriate strategy of combining information across dimensions (rather than attending to only one dimension or just guessing). Again, the GRT does not test for any differences other than attention to dimensions, and any other potential differences in the categorization process were undetected.

SUSTAIN

SUSTAIN is a connectionist prototype model that accounts for both how new items are categorized and how new categories are learned (Love et al., 2004). Instead of a typical prototype model with just one prototype for each category (Hampton, 1993), categories in SUSTAIN are represented by a set of prototypes called subgroups. Subgroups are averages across a subset of previously seen category members, and these subgroups help to represent categories in which there are disparate clusters of entities. For example, spoons are usually either small and metal or large and wooden, but not somewhere in between the two. A prototype model with just one representation of spoons would have trouble representing this structure, so in SUSTAIN spoons might be best represented by two clusters. One averages the features of previously seen teaspoons and soup spoons, and another averages the features of previously seen wooden spoons.

In the categorization process as described by SUSTAIN the similarity between a to-be-categorized item and a subgroup determines the activation of that subgroup. The subgroup with the most activation determines the outcome of the categorization decision. In SUSTAIN and other similarity-based models of categorization, similarity is a mathematical measure of how close the features of a to-be-categorized item are to those of a given category representation (be it an exemplar or a prototype). For SUSTAIN, similarity is determined by degree of match between the to-be-categorized item and the

subgroup's typical feature values, weighted by how important each dimension is for the purposes of categorization. When a dimension is unimportant for identifying a subgroup, any feature value along that dimension adds a little bit to the similarity calculation. For example, the color of a spoon bears little importance to the small spoon cluster, and therefore an item of any color will provide a little bit of activation. For a dimension that is important to category membership, a feature that is located in the correct part of the dimension space adds a lot to the similarity calculation, while a feature located in the wrong part of the dimension space does not add anything at all. For example, having a bowl at its top would provide a lot of activation to both spoon clusters. If a to-be-categorized item is sufficiently similar to one subgroup, it is assumed to have the category membership of that group and the item's features are averaged into the subgroup representation. If it is not sufficiently similar to any subgroup, a new subgroup gets created with the new item's features and category membership.

Like the connectionist prototype models previously discussed, SUSTAIN would have problems accounting for the effects of approach and avoidance on categorization observed by Rock and Janoff-Bulman (2010). First, SUSTAIN would have a hard time accounting for an ephemeral change in dimension saliency since SUSTAIN is a connectionist model that learns the importance of a given dimension over time and there is no mechanism for those values to shift suddenly. This is especially problematic since all context effects that have so far been located within the categorization process have been identified as occurring at the level changing attention to dimensions. Second, it does not have a process for accounting for how similarities or dissimilarities are preferred. Finally, the process of averaging across items to create subgroups obscures the individual information of each previously seen item. If attending to similarities suddenly makes a specific exemplar salient and this impacts categorization decisions, SUSTAIN could only handle this if it preserved that specific exemplar within its own subgroup. For instance, imagine being asked if a honey-dipper was a spoon. It is unlike the average representation

for small spoons, but it may be moderately similar to a previously seen hand-carved wooden spoon. A categorizer attending to commonalities would likely call this spoon to mind while a categorizer attending to dissimilarities would not. SUSTAIN would be unlikely to capture this difference.

GCM

As described in the introduction, the GCM is an exemplar-based model of category application (Ashby & Maddox, 1993; Nosofsky, 1984; Nosofsky, 1986). To categorize an item, the GCM calculates that item's similarity to all previously seen exemplars. The GCM predicts that the item has the probability of being placed in a given category relative to its similarity to exemplars of that category. The GCM has a few aspects that make it attractive for identifying the between-group differences observed by Rock and Janoff-Bulman (2010): it explicitly weights the impact of each dimension of a given item on the similarity calculation, and it has a parameter that scales the impact of moderately similar items.

The GCM has a weight for each dimension that is factored into the categorization decision. These dimension weights allow a dimension to have more weight in certain circumstances and less in others. To return to our feuding Dead Heads example, Categorizer 1 would put all weight on the “shows attended” dimension, Categorizer 2 would put all weight on the “concerts taped” dimension and Categorizer 3 would split weight between both dimensions. Note that these weights are not an inherent property of the exemplars but instead they filter the exemplar at the time of the categorization decision. That is, the dimensions of an exemplar that are important for its category membership are not fixed. This means that, unlike SUSTAIN, the GCM can account for the effects of ephemeral contexts like approach and avoidance on the spread of attention across dimensions (Förster et al., 2006; P. A. Gable & Harmon-Jones, 2008). A person could have attended many Grateful Dead concerts and therefore be similar to other Dead

Heads for Categorizer 1 and also have many concerts on tape and therefore be similar to other Dead Heads for Categorizer 2.

In addition to dimension weights, the GCM has a similarity scaling parameter that converts the distance between a to-be-categorized item and an exemplar to similarity. Similarity decreases exponentially as the distance between that exemplar and that to-be-categorized item increases, and the scaling parameter controls the rate of that decrease. The larger this parameter is, the closer an exemplar needs to be to the to-be-categorized item before it influences the categorization decision. As will be shown, the distance metric in the GCM moves an exemplar farther from a to-be-categorized item when they have non-identical feature values along one dimension but does not move them closer together when they have identical feature values along a different dimension. This means that the similarity scaling parameter scales the effect of dissimilarities on the categorization decision and makes it able to identify the kinds of differences in attention to similarities versus dissimilarities that Forster (2009) identified.

Finally, there is a response scaling parameter that relates certainty in category membership to consistency of categorization response (Ashby & Maddox, 1993). If this value of this parameter is low and people are uncertain about a category membership, their categorization decisions will vacillate. Some days they will say, “yes, the tie-die wearing penthouse dweller is a Dead Head,” and some days they will say “no.” If the value of this parameter is high, and a person is uncertain about category membership, he or she will consistently give whatever answer is most probable.

The GCM has a few drawbacks for identifying the source of context differences. First, the GCM assumes that a to-be-categorized item gets compared to all known exemplars. Depending on how “exemplar” is interpreted, this could mean all memories for every object ever seen. A person simply could not compare an item to all known exemplars when trying to place that item into a natural language category with many members that have each been seen many times. Second, there is debate as to how the

response scaling parameter relates to the categorization process (e.g., Navarro, 2007; Smith & Minda, 2002). While so far it has been discussed as if it relates sum similarity to categorization decisions, it could also be interpreted as directly impacting how similarity is calculated or how exemplars are sampled for the sake of comparison.

The GCM (in different incarnations) has been applied to between-group differences in categorization at least twice. A between-group parameter difference in the GCM was used as evidence against a claim of multiple memory systems, one that is damaged in amnesiacs and one that is intact (Nosofsky & Zaki, 1998). This claim stemmed from the finding that amnesiacs performed at chance on a recognition memory task but performed significantly better at a categorization task that used the same stimuli. A modified version of the GCM was used to predict the categorization performance of amnesiac and normal participants by allowing the similarity scaling parameter to vary between groups. The model did a good job of accounting for between-group differences in categorization when the amnesiac group was given a small similarity scaling parameter, indicating that they were unable to differentiate between similar exemplars in memory. This difference might not impair categorization performance much since category members look more alike than non-category members, but could greatly impair recognition memory.

Additionally, the GCM was used to show the source of between-group differences for participants placed into a happy or sad mood using a traditional, family resemblance set of categorization stimuli (Zivot, Cohen, & Kapucu, 2012). Zivot, Cohen, and Kapucu fit the GCM to each participant's categorization decisions, and found that participants in a happy mood were more likely than participants in a sad mood to spread their attention across multiple dimensions. Happy participants, however, were not as successful as sad participants in identifying the dimensions that allowed them to most effectively distinguish the two categories from each other.

EBRW

The EBRW is the final model considered. It is a random-walk model with an underlying architecture based on the GCM without the response scaling parameter (Nosofsky & Palmeri, 1997). In a random walk model of a categorization decision, information accumulates over time. When a sufficient amount of information is accrued, a decision is made. In the EBRW considered here, “in the category” and “out of the category” represent decisions. Information is accrued in the model by selecting an exemplar and incrementing the category decision a step in the direction of an “in the category” or an “out of the category” decision. The probability that a given exemplar is selected is determined by its similarity to the to-be-categorized item.

For illustration, a mock process of deciding if a honey-dipper is a Spoon is given in Figure 2. At the start of the process, the categorizer has zero knowledge. First, a spoon exemplar gets brought to mind, incrementing the decision process one step towards the “in the category” boundary. Next, a fork exemplar is sampled, incrementing the decision




Figure 2: Example of the random walk process in the EBRW

process one step towards the “out of the category” boundary. Finally, two more spoons are sampled and the “In the category” boundary is reached. At this point, the model predicts that the categorizer makes an “in the category” decision in favor of a honey dipper being a Spoon.

The EBRW shares dimension weighting and similarity scaling parameters with the GCM. Instead of a parameter representing consistency of response, it has two parameters representing the locations of the “In-the-category” and the “Out-of-the-category” boundaries. When both are far from the zero information point, responses will be highly consistent and when both are close to the zero information point, responses will be highly inconsistent. When one boundary is closer than the other to the zero information point, there will be a bias for giving that response relative to the other.

The EBRW addresses some of the issues in using the GCM for modeling the impact of context on categorization decisions. First, it does not represent the categorization decision as comparing a to-be-categorized item to all known exemplars. Instead, the to-be-categorized item is compared to only enough exemplars for a categorization decision to be made. The model calculates the probability that a given exemplar will be influential in the categorization decision, which can be an interesting source of between-group differences in-and-of itself. Second, the EBRW’s boundary parameters allow it to account for an inclusivity or exclusivity bias that is directly built into the categorization process. If participants exhibit an overall bias to put items in a category (or exclude items from a category) it can be directly observed as a change in the boundary parameters.

The added complexity provided by the EBRW comes at a price: the model requires additional data. In order to determine category boundaries (that is, how much information is required for each categorization decision) the EBRW models response times (RT) as well as categorization decisions. This places methodological constraints on

using the EBRW and means that it cannot be used to model categorization data collected without concurrent RT.

Modeling Natural Language Categorization

In order to identify the between-group differences observed by Rock and Janoff-Bulman (2010), any model selected will have to be able to account for natural language categorization decisions. While a few attempts have been made to model context effects on categorization with artificial stimuli, this has not been done for natural language categories. In fact, to date, there have been only a handful of the attempts of applying models of categorization decisions to natural language categories (Verbeemen, Vanpaemel, Pattyn, Storms, & Verguts, 2007; Verheyen, Hampton, & Storms, 2010; Voorspoels et al., 2008) plus one attempt at modeling fruit and vegetable categorization decisions (Smits, Storms, Rosseel, & De Boeck, 2002).

Applying similarity-based models of categorization to natural language categories requires that two major hurdles be overcome. First, similarity-based models require a similarity space that relates items to each other. Second, exemplar-based models require a set of items of known category membership to populate their exemplar space. Researchers have dealt with the similarity problem either by generating a feature applicability matrix (Smits et al., 2002; Verbeemen et al., 2007) or by using participants' similarity judgments (Voorspoels et al., 2008). A feature applicability matrix is created by having one group of participants generate a list of potential features that category members could have and then having a second group determine if each feature is applicable to each potential category member. The resulting matrix is subjected to multi-dimensional scaling (MDS) as a data-reduction technique to generate a tractable set of dimensions. As an example of MDS, imagine every city in the United States was represented by a matrix of distances from every other city. MDS would convert these distances into the set of coordinates for each city that best preserves the distance relationship between cities. For items of known category membership, researchers

assumed that item familiarity was equivalent to category knowledge, and either used familiar items to predict the categorization judgments of unfamiliar items (Smits et al., 2002; Verbeemen et al., 2007) or just used typicality judgments within a category as a dependent measure and assumed that fuzzy category membership would not be problematic (Voorspoels et al., 2008).

A few concerns regarding these methods will be mentioned now and more fully addressed in the Method section of modeling Rock and Janoff-Bulman's (2010) data. First, feature listing causes both theoretical and practical problems. People have been shown to list different features based on context (Wu & Barsalou, 2009), and feature lists are a binary judgment (i.e., present vs. absent) whereas many features are continuous or fuzzy. Methods for determining category membership for the purposes of modeling categorization decisions are also a cause for concern. Category membership determined a priori by the experimenters depends upon the experimenter's judgment for the answer to the very question being asked of participants in the experiment (Smits et al., 2002; Verbeemen et al., 2007). If questions of category membership were clear, there would not be fuzzy categories or court cases about the category membership of tomatoes. Attempts to circumvent the problem by just looking at typicality as measured by within-category similarity (Voorspoels et al., 2008) are also problematic because they ignore evidence that typicality is related to both increased similarity to In-the-category members and decreased similarity to "out of the category" members (Rosch & Mervis, 1975; Verheyen, De Deyne, Dry, & Storms, 2011).

In spite of these concerns, the studies of natural language categorization cited above have been a good first step, showing promising methods and addressing issues of concern to categorization as a field. As a proof-of-concept, Smits et al. (2002) showed that models of categorization decisions can be applied to categories learned outside the laboratory and that people's knowledge of category members are good predictors of their categorization behavior when encountering an unknown item. Both

Smits et al. and Voorspoels et al. (2008) tested whether an exemplar model (the GCM) or prototype model provided a better fit of categorization responses and found that the GCM accounted for categorization decisions better. Verbeemen et al. (2007) used natural language categories to test the Varying Abstraction Framework (Vanpaemel & Storms, 2008), a subgrouping model of categories that exists as an intermediary between prototype and exemplar theories. Previous findings in favor of a subgrouped representation are strengthened by their success modeling real-world items. In an ideal world, however, these subgroups would be intuitively sensible (e.g., there is no identifiable reason why nectarines, oranges, tangerines, plums, and cherries and all placed into a single group).

The GCM Defined

As discussed above, the GCM is promising model for accounting for the between-group differences observed by Rock and Janoff-Bulman's (2010). First, the parameters of the GCM capture a number of the components of categorization that prior research has shown can be affected by approach and avoidance mode. Second, the GCM has been previously been used to model natural language categories. This makes it a good starting point for identifying the stage in categorization where avoidance mode affects Conservatives' categorization decisions (Rock & Janoff-Bulman, 2010). For the purposes of this dissertation, the GCM will be used to model between-group differences when RT is not available. This section will outline the formal properties of the model.

Formally defined, the first step in the categorization process according to the GCM is to calculate the distance between the to-be-categorized item and all exemplars within a similarity space. The formula for calculating the distance d between to-be-categorized item i and exemplar j is,

$$d_{ij} = \sqrt{\left| \sum_m w_m |x_{im} - x_{jm}|^2 \right|} \quad (1)$$

where m is the dimension being evaluated, w_m is the weight placed on that dimension, and x_{jm} is the value of exemplar j along dimension m . In this model, all dimension weights must be greater than or equal to zero, and must sum to one. These dimension weights allow one dimension to be favored over another for the purposes of categorization. If, for some reason, tie-color was important for telling Republicans from Democrats but shoe color was not, weight would be placed on the tie dimension ($w_{\text{tie}} = 1$) but not on the shoe dimension ($w_{\text{shoe}} = 0$).

The next step in the categorization process is to convert distance to similarity (η). The η_{ij} between to-be-categorized item i and exemplar j is determined by the exponential decay function defined by Shepard,

$$\eta_{ij} = \exp(-c \cdot d_{ij}) \quad (2)$$

where c is the similarity scaling parameter. With a high c value, only the only exemplars considered similar to the item are the ones that are very close in similarity space.

The probability of an “In-the-category judgment” (p_i) can be found by taking its similarity to exemplars in category A (S_{iA}),

$$S_{iA} = \sum_{j \in A} \eta_{ij} \quad (3)$$

as well as its similarity to exemplars not in category A (S_{iB}),

$$S_{iB} = \sum_{j \in B} \eta_{ij} \quad (4)$$

and dividing S_{iA} by the sum similarity to all exemplars,

$$p_i = \frac{S_{iA}}{S_{iA} + S_{iB}} \quad (5)$$

The process for determining the probability of giving a “not-In-the-category” response (q_i) is similar,

$$q_i = \frac{S_{iB}}{S_{iA} + S_{iB}} \quad (6)$$

The EBRW Defined

As noted in the section Modeling the Categorization Process, the EBRW’s properties also make it a good candidate for accounting for the between-group differences observed by Rock and Janoff-Bulman’s (2010). First, its parameters are even more

analogous than the GCM to the aspects of cognition shown to be affected by approach and avoidance modes. Furthermore, previous efforts made at fitting the GCM to natural language categories can be easily translated to the EBRW. For the purposes of this dissertation, the EBRW will be used to model between-group differences when RT is available. This section will outline the formal properties of the model.

The EBRW is a random walk model of categorization decisions, where two categories are represented as absorbing states at the boundaries of the random walk space. According to the EBRW, the categorization process occurs by sampling exemplars from memory. The probability that an exemplar is sampled is relative to its similarity to the to-be-categorized item. Each exemplar sampled moves the random walk state towards the category boundary that the exemplar belongs to. Reaching one of the absorbing boundaries represents the end of the categorization decision.

At its core, the EBRW repurposes the formulas from the GCM. According to the EBRW, p_i (the probability of placing item i into category A) becomes the probability that an exemplar from category A will be sampled by memory and therefore the probability that a step will be taken towards the category A boundary (Figure 2). Likewise, q_i becomes the probability that a step will be taken towards the category B boundary. The probability that an item i will be put into category A is represented by

$$P(A | i) = \frac{1 - (q_i / p_i)^B}{1 - (q_i / p_i)^{A+B}} \quad (7)$$

The parameter A in this equation represents the “In-the-category” boundary – its distance above the zero-information point. Likewise, parameter B in this equation represents the “Out-of-the-category” boundary – its distance below the zero information point. As one becomes larger than the other, the model will predict a bias to respond with the category that has the closer boundary. As both boundaries get farther from the zero-information point, the model predicts less probability matching and more consistent categorization responses.

The EBRW also uses the sum similarity (Equation 3) to predict response time to each stimulus. Specifically, it predicts that the time to take any step in the random walk process ($E(T_{\text{step}})$) for item i is given by,

$$E(T_{\text{step}} | i) = \alpha + 1(S_{iA} + S_{iB}), \quad (8)$$

where α is a constant that represents the time for cognitive tasks that are independent of the similarity calculation and therefore are common across all items, such as time to visually encode the jellyfish or press a button once a categorization decision has been reached. Note that as similarity to either category increases, the time required to retrieve an exemplar and make a step in the random walk process decreases.

The EBRW predicts that the total number of steps required during the random walk process ($E(N)$) for a given item i is a function of both the probability of placing item i in the category or not (Equations 4 and 5), as well as the distance of the decision boundaries A and B from the zero-information point (for in and Out-of-the-category, respectively):

$$E(N | i) = \frac{B}{q_i - p_i} - \frac{A + B}{q_i - p_i} \left[\frac{1 - (q_i / p_i)^B}{1 - (q_i / p_i)^{A+B}} \right] \quad (9)$$

The predicted amount of time it should take to categorize an item can be found by multiplying Equation 8 by Equation 9.

Locating the Effects of Mode in the EBRW

To better understand how these parameters can account for the effects of motivation on cognition, an example category decision will be modeled with the EBRW (Figure 3). Four exemplars are known: two from category A and two from category B, while the to-be-categorized item is represented by ‘?’’. Each exemplar has two dimensions that can take a value from 1 to 5, plotted along axes X and Y. According to the EBRW, when attention is evenly divided across both dimensions (i.e., $w_1 = 0.5$, $w_2 = 0.5$), the similarity scaling parameter is small (i.e., $c = 1$), and both boundaries are close and equidistant from the zero-information boundary (i.e., $A = 1$, $B = 1$), then there is a 33% chance of putting item “?” into category A, and a 66% chance of putting it into




Figure 3: Example of a categorization decision as described by the EBRW

category B. This is because the to-be-categorized item is most similar to the category B exemplar (5,2), followed by the category A exemplar (3,1).

To see how these parameters impact categorization decisions and capture the effects of motivation on cognition, imagine putting all the attention weight on dimension x (i.e., $w_1 = 1, w_2 = 0$) while holding all other parameters the same. This would collapse the exemplars across the y dimension. Now, the to-be-categorized item is identical to both category B exemplars but it would remain equally dissimilar to the category A exemplars. The EBRW now predicts a 7% probability “?” will be put into category A. With all the weight on dimension y , however, the to-be-categorized item becomes identical to both A exemplars. Now the EBRW predicts an 82% probability that “?” will be put into category A. Note that the two values are not reciprocal. This is because the to-be-categorized item is more similar to exemplar B(5,2) when the x dimension is collapsed than it is to exemplar A(3,1) with the y dimension collapsed. The EBRW’s ability to change dimension weights makes it a good model to test the theory that motivational states affect

the way attention is distributed during categorization (e.g., Förster et al., 2006; P. A. Gable & Harmon-Jones, 2008)

Now imagine that attention is equally distributed across dimensions but the similarity scaling parameter (c) is increased. As c gets larger, the probability that the to-be-categorized item will be placed into category A decreases exponentially from .20 when $c = 2$, to .06 when $c = 4$, to essentially zero when $c = 7$. This is because when the similarity scaling parameter is low, items that are only moderately similar (such as A(3,1)) can still have an impact on the categorization decision. When c is high, however, only the most similar items can influence categorization decisions (i.e., B(5,2)). In this way, the EBRW can test the theory that approach and avoidance impacts whether people attend to what makes things similar or what makes them different. A low c parameter reflects a categorizer that is looking for similarities and is including information from exemplars that are only moderately similar. A high c parameter reflects a categorizer that is looking for dissimilarities and is only including information from the most similar exemplars, pushing away moderately typical exemplars. This makes the EBRW a good model for testing the theory that motivational state changes whether a person looks for similarities or dissimilarities (e.g., Förster, 2009).

Changing the boundary parameters, A and B, can have two different types of effects depending on whether they move equidistantly from the zero-information point or one moves more than the other. As A and B equally increase in distance from the zero-information point, categorization becomes more deterministic. A categorizer with boundaries far apart from each other would be highly consistent—more likely to repeatedly put the same item into the same. In this example, the to-be-categorized item has twice as much similarity to exemplars from category B than category A. When boundaries A and B both equal 1, the model predicts the categorizer will say B 66% of the time. As the boundaries move farther apart, the model predicts the categorizer will say

B more and more frequently: 89% of the time when A and B equal 3 and 97% of the time when A and B equal 5.

When one boundary increases and the other remains the same, this creates a bias to respond to the category with the closer boundary. A categorizer with a bias to respond “In-the-category” would be a more inclusive categorizer and a categorizer with a bias to respond “Out-of-the-category” would be a more exclusive categorizer. In the current example, as A remains 1 but B moves farther away, the probability of saying A increases from 33% when $B = 1$, to 46% when $B = 3$ and plateaus at 49% when $B = 5$ and higher. In sum, between-group differences in boundary locations in the EBRW model can indicate whether participants in one motivation condition are changing their decision biases (Markman et al., 2005) or are making more consistent categorization decisions.

CHAPTER 3

OVERVIEW OF EXPERIMENTS

Rock and Janoff-Bulman (2010) showed that Conservatives in avoidance mode were more exclusive categorizers of natural language category members than Conservatives in approach mode. They did not find a difference between Liberals in approach and avoidance mode. The purpose of this dissertation is to use two exemplar-based models of categorization, the GCM and EBRW, to describe these between-group differences.

Experiment 1 fits the EBRW to the categorization decisions of Conservative and Liberal participants in approach and avoidance modes using artificially generated stimuli. The stimuli in Experiment 1 have clearly defined dimensions that can be accounted for by mathematical models of categorization. These artificial stimuli are typical of the categorization experiments to which the EBRW has previously been applied (Nosofsky & Palmeri, 1997; Nosofsky & Alfonso-Reese, 1999). The purpose of this experiment is to use the EBRW to identify the components in the categorization process where this between-group difference occurs.

Experiment 2 fits the GCM to data from Rock and Janoff-Bulman's (2010) Experiment 1. Exemplar-based models have been fit to natural language categorization decisions (Verbeemen et al., 2007; Voorspoels et al., 2008). There has not been consistency across prior studies in which the GCM has been fit to natural language categories, and some methods are not amenable to identifying between-group categorization differences. This experiment has two purposes: 1) to provide a testing ground for generating information required by both the GCM and the EBRW to fit between-group differences in natural language categorization decisions; and 2) to identify differences in the components of categorization that underlie the between-group differences observed by Rock and Janoff-Bulman.

Experiment 3 fit the EBRW to data that replicated Rock and Janoff-Bulman's (2010) Experiment 1 for two categories and extended it by collecting RTs and adding stimuli to each of the two categories. Rock and Janoff-Bulman collected data on categorization decisions for 12 items per category and do not record RT. Parameters of the EBRW may provide a better analog than the GCM for differences in cognition observed between people in approach and avoidance mode, but the EBRW cannot be fit to Rock and Janoff-Bulman's data because they did not record RT. The purpose of Experiment 3 is to fit the EBRW to natural language categories in order to locate the source of differences between Liberals and Conservatives in approach and avoidance mode in categorization decisions.

CHAPTER 4

EXPERIMENT 1

Rock and Janoff-Bulman (2010) found that Conservatives in avoidance mode were more exclusive categorizers of moderately typical items than Conservatives in approach mode, but did not find this difference for Liberals. Experiment 1 uses the EBRW to show the components in the categorization process in which this between-group difference occurs by replicating Rock and Janoff-Bulman's findings with artificial stimuli. The EBRW has been shown to successfully model categorization decisions using these types of stimuli (e.g., Nosofsky & Palmeri, 1997; Nosofsky & Alfonso-Reese, 1999).

The EBRW is a promising model for locating these effects because aspects of cognition that approach and avoidance mode have been shown to impact have analogs in the model's parameters. Approach and avoidance mode have been shown to affect the spread of attention across dimensions (Friedman & Förster, 2005; P. A. Gable & Harmon-Jones, 2008). The EBRW has parameters representing the amount of attention given to each dimension of the to-be-categorized item. Approach and avoidance mode impact how much attention is given to dimensions along which items are similar versus dimensions along which items are different (Förster, 2009). The EBRW has a similarity scaling parameter that scales the effect of dissimilarities on a categorization decision. Approach and avoidance mode can impact where decision criteria are established (Markman et al., 2005). The EBRW has random walk boundaries *A* and *B*, which can change independently representing a bias to respond "In-the-category" or "Out-of-the-category." Additionally, as *A* and *B* move farther away from each other, categorization decisions become more consistent.

Experiment 1 uses artificial stimuli: computer drawn jellyfish that vary along four dimensions. Similar stimuli have been used in many efforts to model categorization judgments (e.g., Kruschke, 1992; Love et al., 2004; Nosofsky, 1984). As discussed

in the Literature Review, these stimuli offer a number of attractive features for operationalizing the stages of the categorization process. First, the stimuli are unknown to the participants, who theoretically do not enter the experiment with preconceived ideas about which features are important and which are not. Second, they have a limited number feature values (e.g., round vs. square heads, color saturation values, line length) along dimensions that are obvious to the categorizer and are amenable to mathematical modeling. Third, these stimuli can have family resemblance structures built into them, making them potential stand-ins for complex real-world categories (e.g., Medin & Schaffer, 1978).

Stimuli typically used in categorization experiments are comprised of dimensions with binary feature values and are often based on the 5/4 category structure created by Medin and Schaffer (1978). In real world items, however, features are rarely just present or absent; they usually exist along a continuum (e.g., height and color) or are ambiguous (e.g., having wings—consider flying squirrels). Furthermore, the use of binary feature values gives rise to concerns that participants encounter so few permutations that they merely memorize which items are associated with which category labels (Blair & Homa, 2003). The use of stimuli with features that vary continuously along a set of dimensions mitigates this concern. The features of the jellyfish used in this experiment vary continuously along four dimensions: 1) color of the internal organs, 2) diameter of the bell, 3) thickness of the bell, and 4) tentacle length.

Method

Participants

Participants were recruited from University of Massachusetts, Amherst psychology subject pool. Students in this pool answered four questions about their political identity as part of a questionnaire administered at the beginning of the semester (see Appendix A). Participants were recruited based on their responses to these questions. The answer to the fourth question was reverse-coded and all four responses were

averaged together to generate a mean political identity score with a range of 1(Liberal) to 7(Conservative). In Spring Semester 2011, all pre-screening respondents (n=1200) had an average political identity score of 3.4 with a standard deviation of 1.0. Students were identified as liberal if their score was one standard deviation or more below the mean or as conservative if their score was one standard deviation or more above the mean.

Students identified as liberal or conservative were contacted via email and invited to participate in either this experiment, identified as “Categorizing Jellyfish,” or in Experiment 3, identified as “Categorizing Everyday Objects.” Recruitment occurred across three semesters. The mean and standard deviation of this political score was essentially unchanged across all three semesters. For the Spring Semester 2011 $\mu = 3.4$ and $\sigma = 1.0$, for the Fall Semester 2011 $\mu = 3.5$ and $\sigma = 1.0$, and for the Spring Semester 2012 $\mu = 3.5$ and $\sigma = 1.0$. The selection criterion from Spring Semester 2011 was used for all three semesters. This selection criterion identified roughly one-third of the students as potential participants each semester. Participants were not allowed to sign up for both studies.

One hundred and twenty seven University of Massachusetts, Amherst students participated in this study (33 males and 106 females). Data from four of these participants were not included in this analysis due to computer error. Of the remaining participants, 50 were identified as political conservatives and 73 were identified as political liberals. Forty participants were placed into the approach condition (20 Liberals, 20 Conservatives), 41 in the neutral condition (23 Liberals, 18 Conservatives), and 42 in the avoidance condition (30 Liberals, 12 Conservatives).

Materials

The stimuli for this experiment consisted of 16 jellyfish, generated in Psychtoolbox (Brainard, 1997; Pelli, 1997). For example jellyfish, see Figure 4. The jellyfish were designed with features that varied continuously along four dimensions: 1) color of the internal organs, 2) diameter of the bell, 3) thickness of the bell, and 4)

tentacle length. The stimuli were generated according to one of two methods, one for category member items and one for non-category items.

The category member jellyfish were generated by adding normally distributed noise to each feature value of the prototype. The prototype had organ colors with an RGB values of [145, 0, 0], a bell diameter of 420×210 pixels, a bell thickness of 6.45 points, and a tentacle length of 25 pixels per segment. (There were nine segments.) Figure 4B shows the prototype jellyfish. Three types of category member jellyfish—highly typical, moderately typical, and atypical—were generated based on a prototype, which was not shown during the experiment. Three jellyfish were generated for each level of typicality, for a total of nine category member jellyfish stimuli. For highly typical category members, the normally distributed noise was set at one standard deviation for each dimension. (See Appendix Table C1 for standard deviation values). For moderately typical category members, normally distributed noise was set at two standard deviations for each dimension. For atypical category members, normally distributed noise was set at three standard deviations for each dimension. These category member jellyfish were considered as belonging to a category since they all had features similar to the prototype and therefore similar to each other. Highly typical jellyfish were most likely to have all their features similar to each other and dissimilar to non-category member jellyfish. Seven non-category member jellyfish were generated by selecting each dimension's feature value at random from a uniform distribution. Figure 4A and Figure 4C present




Figure 4: Example jellyfish stimuli. A) Minimum non-category jellyfish values, B) Prototype jellyfish values, C) Maximum non-category jellyfish values (Participants did not see any of these jellyfish during the experiment.)

the minimum and maximum values for the non-category member jellyfish, respectively. Feature values for each of the 16 stimuli are given in Appendix Table C2. All participants saw the same 16 jellyfish.

Procedure

Upon arriving at the experiment, participants gave their informed consent to participate and were brought to a room containing an individual computer. Participants received all instructions for the experiment from this computer. After being welcomed to the study, they were told that a new species of jellyfish had been discovered (*Aurelia diatribi*) and their task was to learn to distinguish category member jellyfish from non-category member jellyfish. During the introduction, they were told the four dimensions essential to categorizing the jellyfish.

In the training stage of the experiment, participants were trained on four category member jellyfish (High 1-3 and Medium 1) and four non-category jellyfish (Non-category 1-4). These jellyfish were selected for training to provide participants with a strong sense of the feature values associated with the category and to allow for as much variability in categorizing atypical items as possible, since participants were not trained on atypical stimuli. In Appendix Table C2, all training stimuli are bolded.

On every trial during training, participants were shown a jellyfish and asked if it belonged to the category. After making a categorization decision, they received feedback; either “Correct” accompanied by a pleasant tone or “Incorrect” accompanied by a beep. A training block consisted of a run through all eight training jellyfish, presented in a random order. Training continued until either a participant correctly categorized two blocks of jellyfish in a row or they completed 15 blocks.

In the second stage, participants were primed for approach mode, avoidance mode, or a neutral condition. This priming was based on the amoral prime reported in Rock and Janoff-Bulman (2010). Participants in all conditions were given five minutes to list up to ten movies. Participants in the approach condition were asked to list movies to

watch for an enjoyable time, those in the neutral condition were asked to list movies they had seen recently, and those in the avoidance condition were asked to list movies to not watch to avoid a bad time. These primes were intended to focus approach participants on what they should do (i.e., movies they should watch) and avoidance participants on what they should not do (i.e., movies they should avoid). The effectiveness of these primes was considered exploratory in Rock and Janoff-Bulman. The full text of the priming instructions is provided in Appendix B.

In the third stage, participants were asked to categorize jellyfish without feedback. In this stage, each block consisted of sixteen jellyfish: the eight that participants had been trained on and eight new jellyfish: two moderately typical, three atypical, and three non-category member jellyfish. On each trial, participants saw a jellyfish, made an “In-the-category” or an “Out-of-the-category” judgment, and were thanked rather than given feedback. Participants completed ten blocks of category judgments and the order of jellyfish presentation was randomized within each block. Therefore, the participants categorized the 16 jellyfish ten times for a total of 160 categorization judgments. This stage took approximately five minutes.

In the fourth stage, participants were asked to list three goals they aspired to achieve and three goals they felt obligated to achieve. The order of these two tasks was counterbalanced across participants. Response times to these questions are taken as a check of successful priming with approach or avoidance mode (Friedman & Förster, 2001). Participants in the approach mode are expected to respond faster to desires while participants in avoidance mode are expected to respond faster to obligations.

Finally, in the fifth stage, participants rated the paired similarity of all sixteen jellyfish in the study. On each trial, participants saw one jellyfish and then pressed a button to see the next. They were asked to decide their similarity on a scale of 1 (most dissimilar) to 9 (most similar). This scale remained on the bottom of the screen at all times, and participants were encouraged to use the whole scale. Pair presentation order

was random and participants were not asked to make similarity judgments about identical jellyfish.

After completing all five stages, participants were debriefed and thanked for their participation. On average, the experiment took participants 30 minutes to complete.

Results

Manipulation check

The measure of interest for the manipulation check was calculated according to Friedman (2001): the amount of time it took participants to start typing each of the three goals they aspired to accomplish and each of the three goals they were required to accomplish. Differences in the time to start typing were tested by a 3(Mode) \times 2(Question Type) mixed model ANOVA, with Question Type as a within-subjects factor. First, there was a main effect of Question Type. Participants were faster to start typing goals that they aspired to achieve ($\bar{x} = 6.35, s = 3.53$) than goals that they were required to achieve ($\bar{x} = 10.04, s = 5.63$), $F(1,116) = 101.96, p < .001$. There was, however, neither a main effect of Mode ($F(2,116) = 0.81, p = .45$) nor an interaction between Mode and Question Type ($F(2,116) = 0.35, p = .71$).

Category learning

A participant was considered to have successfully learned the category if they either got two blocks correct in a row or got seven or more out of eight correct on the last block of training. By this measure, 91 out of the 123 participants successfully learned the category. Category learning did not differ either by Mode, $\chi^2(2) = 0.22, p = .90$, or Political Identity $\chi^2(1) = 0.18, p = .67$. Only participants who learned the category were included in further categorization performance and modeling analyses.

Categorization performance

Rock and Janoff-Bulman (2010) found that Conservatives primed with avoidance mode were more exclusive when categorizing moderately typical category members as compared to Conservatives primed with approach mode. The category inclusivity for

Liberals, however, was the same regardless of whether they were primed with approach or avoidance mode. Since it is the goal of this study to identify the stage in categorization that approach and avoidance motivation effects, it is first necessary to show that there was a difference in the categorization performance of these two groups in this study as well.

In order to identify between-group differences in categorization performance, the percentage of jellyfish included in the category by each participant was averaged across highly typical members, moderately typical members, atypical members, and non-members, resulting in four category-inclusion measures for each participant. As can be seen in Table 1, category inclusion is high for highly typical and moderately typical category members, less for atypical members, and low for non-category members. These results are broken down by gender in Appendix Table F1 and Appendix Table F2.

To normalize the distribution, categorization inclusion percentages were arcsine transformed and then submitted to a 4(Typicality) \times 3(Mode) \times 2(Political Identity) mixed model ANOVA, with Typicality as a within-subjects factor. As expected, there was a main effect of Typicality, $F(3,255) = 534.0, p < .001$. Additionally, there was an interaction between Mode and Political Identity, $F(2,85) = 4.2, p < .05$. This means that, averaged across category typicality, there was a difference between Liberals and Conservatives in how mode affected category inclusion. The three-way interaction between Mode Manipulation, Political Identity, and Typicality that would have directly

Table 1
Average %yes Judgments (and SE) by Manipulation and Stimulus Type

Group	High	Medium	Atypical	Non-category	Total
Liberal					
Approach	.88(.02)	.91(.02)	.76(.04)	.21(.03)	.69(.02)
Neutral	.83(.03)	.83(.04)	.68(.04)	.16(.04)	.62(.03)
Avoidance	.89(.03)	.93(.02)	.75(.03)	.21(.02)	.69(.02)
Conservative					
Approach	.83(.03)	.78(.05)	.68(.03)	.19(.03)	.62(.03)
Neutral	.90(.05)	.90(.05)	.70(.05)	.18(.02)	.67(.03)
Avoidance	.91(.02)	.93(.02)	.70(.04)	.13(.02)	.67(.02)
Total	.87(.01)	.88(.02)	.71(.02)	.19(.01)	

replicated Rock and Janoff-Bulman's (2010) results was not significant, $F(6,255) = 1.2$, $p = .308$.

To investigate the interaction between Mode and Political Identity, two different analyses were performed. First, the proportion of "In-the-category" decisions was averaged across all four typicality levels (Table 1), creating an overall measure of category inclusion. Since the result of interest is how approach and avoidance modes affect categorization decisions, this measure was subjected to two follow-up analyses, one testing the effect of approach mode and the other testing the effects of avoidance mode. For participants in approach mode, Liberals were more inclusive categorizers than Conservatives (Table 1), $F(1,28) = 4.4$, $p < .05$. For participants in avoidance mode, there was no significant difference between Liberal and Conservative participants, $F(1,28) = .885$, $p = .355$.

Additionally, a signal detection analysis was performed on these data. Signal detection decomposes hit rates and false alarm rates into a measure of participants abilities to distinguish category members from non-category members (measured in d'), as well as their criterion for placing an item in the category versus out of the category (measured in c_{SD}). A c_{SD} of 0 indicates that a categorizer is optimizing their judgments to maximize percent correct, a negative c indicates that the categorizer is including more to-be-categorized items than optimal, and a positive c indicates that the categorizer is excluding more to-be-categorized items than optimal. Hit rates were measured as the percentage of all category members included in the category and false alarms were measured as the percentage of all non-category members included in the category. Hit rates at ceiling were transformed by subtracting one included category member and false alarm rates at floor were transformed by adding one included non-category member. The average results for these analyses are reported in Table 2.

Both d' and c_{SD} were submitted to a 3 (Mode) \times 2 (Political Identity) ANOVA. For d' , there were no significant between group differences. The ability to distinguish

Table 2
Average Parameter Values (and SE) for Signal Detection Analysis

Group	d'	c_{SD}
Liberal		
Approach	1.95(0.12)	-0.10(0.07)
Neutral	2.08(0.15)	0.12(0.11)
Avoidance	2.08(0.12)	-0.18(0.07)
Conservative		
Approach	1.74(0.18)	0.08(0.08)
Neutral	2.08(0.18)	-0.09(0.10)
Avoidance	2.27(0.09)	0.08(.09)

category members from non-category members did not differ between Mode ($F(2,85) = 2.30, p = .11$), between Political Identity ($F(1,85) = 0.01, p = .94$), nor was there an interaction between the two ($F(2,85) = 0.81, p = .45$). For c_{SD} , while there was no main effect of Mode ($F(2,85) = 0.24, p = .79$) or Political Identity ($F(1,85) = 0.97, p = .33$), the interaction between the two was significant ($F(2,85) = 3.75, p < .05$). Follow up analyses of this interaction show that there are no significant differences between participants Liberals and Conservatives in approach mode, $F(1,28) = 2.50, p = .13$. Liberals in avoidance mode, however, have a more inclusive criterion than Conservatives in avoidance mode, $F(1,28) = 4.32, p < .05$.

Modeling

Participants' categorization decisions and response times were fit with the EBRW. As outlined in the Literature Review, the EBRW combines an item's similarity to the exemplars in the category (S_{iA}) with its similarity to exemplars out of the category (S_{iB}) to predict both the probability that the item will be categorized as "In-the-category," as defined by Equation 7, as well as the RT to that item, as defined by Equations 8 and 9. Since the EBRW predicts RT in arbitrary units, predicted RT is scaled to milliseconds by a linear regression, with slope k and intercept μ . For each model, the EBRW was fit to four different data sets: categorization responses and response times averaged within Liberals and Conservatives in approach and avoidance mode.¹

1 In Experiment 1, the EBRW can be fit to individual participants' data as well as

For the first step in the modeling process, a fully constrained version of the EBRW was fit to all four data sets. This model found the c , the set of w parameters, the A and B boundaries, and α that best minimized the weighted sum of squared deviations ($WSSD$) between the observed and predicted proportion of “in-the-category” judgments and observed and predicted average response times for each of the 16 stimuli. The $WSSD$ method was proposed by Nosofsky and Stanton (2005) as a way to combine these two error sources, which are measured on different scales. $WSSD$ weights the error of each prediction by the inverse of the data point’s squared standard error. Therefore, category decisions and RT are standardized to a similar scale and data points with smaller variability play a larger role in the fit measure. The fully constrained model was run 100 times with random starting points to guard against local minima, and the parameters from the best fitting model were kept. Next, this model was run again 100 times with the best fitting parameters as its starting point, with each parameter permuted by normally distributed noise ($N(0,1)$).

The next step in model fitting was to allow parameters to vary between groups. The data were fit with models that allowed either c , all w weights, A , B or both A and B to vary between groups. Parameters were allowed to vary one of three ways: between Liberals and Conservatives participants, participants in approach mode and avoidance mode, and Conservatives in avoidance mode compared to all others. This last model was included because of the interaction found by Rock and Janoff-Bulman (2010). Each model was fit 100 times using the best fitting parameters of the fully constrained model plus normally distributed noise as starting points. The predictions of the best fitting models can be found in Appendix Table F4.

Since the EBRW has no known likelihood fit measure (Nosofsky & Stanton, 2005), models in Experiment 1 were compared with cross-validation (Browne, 2000). averaged data. Although this method is recommended when possible (Maddox, 1999), to do so for Experiment 1 would mean that each data point being modeled is the average of only ten trials. The EBRW was fit to individual participant’s data and statistical tests of the parameters are reported Appendix Table F3.

Cross-validation fits the model to one set of data and tests it on another. It is based on the assumption that a model that is too complex will adjust its parameters to fit noise in the data and do a poor job of predicting new data. Cross-validation was implemented by cycling through each to-be-categorized item, fitting each of the models with that item's categorization and RT data withheld, and then using the resulting parameters to predict the withheld data. Each model's cross-validation fit measure is the average of its *WSSD* across all predicted data points, and this value can be found in the Appendix Table F4.

The best fitting model allowed the *c* parameter to vary between Conservatives in avoidance mode and all other participants. Figure 5 shows the averaged category inclusion rates and model predictions and Figure 6 shows the averaged RT and model predictions. These data were best fit by a model that allowed Conservatives in avoidance mode to have a larger *c* parameter ($c_{\text{ConAvo}} = 1.20$) than other participants ($c_{\text{Other}} = 0.58$). This indicates that, for Conservatives in avoidance mode, the functional relationship between psychological distance and psychological similarity dropped off more steeply than for other participants.

Discussion

Liberal and Conservative participants in Experiment 1 were primed with approach or avoidance mode and then made categorization decisions about a newly learned category. Analyses of the categorization performance indicate that, overall, Liberals were more inclusive categorizers than Conservatives. ANOVA analyses show that in approach mode they are more inclusive overall and signal detection analyses show that in avoidance mode, Liberals have a more inclusive criterion than Conservatives. When the EBRW model of categorization was fit to individual participant's categorization decisions, analyses indicate that Conservatives in avoidance mode had a larger *c* parameter. These results provide evidence that exemplar-based models in general, and the EBRW specifically, can be used to account for between-group differences in categorization decisions. While the manipulation check did not find significant




Figure 5: Average and model predicted percent “in-category” judgments by stimulus




Figure 6: Average and model predicted RT by stimulus

differences between-groups, inferential statistics of the categorization performance indicate that the manipulation may have been effective.

In previous applications of the GCM (e.g., Nosofsky, 1987; Nosofsky & Zaki, 1998) and the EBRW (Nosofsky & Palmeri, 1997), the scaling parameter is taken as a measure of memory - how well participants can distinguish one exemplar from another. Nosofsky and Palmeri showed that the categorization and memory performance differences between normal participants and amnesiacs could be modeled with just a change in the scaling parameter, indicating that the amnesiacs' deficit came from a reduced ability to distinguish between exemplars. Nosofsky (1987) showed that when participants learned a category over many learning trials across a number of sessions, the scaling parameter increased across sessions, indicating an improved ability to distinguish between exemplars.

It is not necessary to conclude, however, that being in avoidance mode increased the Conservative participants' ability to distinguish between exemplars but did not do so for Liberal participants. Another interpretation is that avoidance mode made Conservative participants less willing to factor moderately similar exemplars into their categorization decisions. For these conservative participants, the to-be-categorized items were put into the same category as exemplars that were highly similar to them. The distance calculation of the EBRW (Equation 1) starts at zero distance between an item and an exemplar and only increases when there are dissimilarities between the two. An increased similarity scaling parameter, therefore, can be interpreted as placing an increased weight on dissimilarities.

There are two theories about the way in which approach and avoidance mode impacts the distribution of attention. Some researchers have found evidence that approach and avoidance modes modulate the focus of attention from the big picture to the fine details (e.g., Friedman & Förster, 2005; P. A. Gable & Harmon-Jones, 2008), while others have found that approach and avoidance modulates how attention is distributed across

dimensions that highlight similarities or differences (Förster, 2009). These theories are not necessarily in competition with each other and parameter differences in the EBRW could have found evidence for both. That said, the results of this study can be taken as evidence for the theory that avoidance focuses people on dissimilarities. This effect, however, was only found for Conservatives. Liberals were likely not as susceptible to the avoidance mode manipulation, since approach mode is found to be dominant in the personality of Liberals (Janoff-Bulman et al., 2008; Rock & Janoff-Bulman, 2010).

The impacts of the mode manipulation were modest in Experiment 1. The category-inclusion ANOVA analysis, the signal detection analysis, and the modeling analysis all found significant differences between different groups within the study. It is possible that the approach and avoidance manipulation did not succeed in significantly manipulating participants' mode or perhaps the task of categorizing previously unseen items washed out much of the experimentally generated context. What is most important for the purposes of this dissertation, however, is that the EBRW was able to capture the source of reliable difference between Liberals and Conservatives.

CHAPTER 5

EXPERIMENT 2

Experiment 1 showed that an exemplar-based model of categorization, the EBRW, can be used to identify between-group differences in components of the categorization process for artificial stimuli. In order to account for the effects of avoidance mode observed in Conservatives by Rock and Janoff-Bulman (2010), exemplar-based models must also be able to identify between-group differences for natural language category stimuli. The GCM, a related exemplar-based model of categorization, has already been applied to natural language categories (Verbeemen et al., 2007; Verheyen et al., 2010; Voorspoels et al., 2008). None of these studies of natural language categorization, however, have used the GCM to identify between-group differences in components of the categorization process.

To show that exemplar-based models can identify components of the categorization process in which natural language categorization differs between categorizers in approach versus avoidance mode, Experiment 2 will fit the GCM to the data collected by Rock and Janoff-Bulman (2010). The GCM is used to model these data because the GCM does not require RTs, which were not collected by Rock and Janoff-Bulman. As discussed in the Literature Review, the GCM compares a to-be-categorized item to exemplars and categorizes it based on the category membership of the exemplars to which it is most similar. The GCM shares a number of parameters with the EBRW (c and the w attention weights). These parameters break down the categorization process in a way that provides analogs with the ways that approach and avoidance mode have been shown to impact the spread of attention.

Fitting the GCM to these data will provide a test for generating information required by both the GCM and the EBRW to fit between-group differences in natural language categorization decisions, such as identifying objective methods of generating a similarity space and supplying a set of exemplars. Previous researchers who have applied

the GCM to natural language categories have based their similarity space on a subjective list of features for each exemplar (e.g., Smits et al., 2002; Storms, De Boeck, & Ruts, 2001; Verbeemen et al., 2007; Voorspoels et al., 2008). Typically, these researchers have one group of participants generate a list of features that could apply to the whole category and have a second group decide whether these features are applicable to each exemplar. The resulting feature matrix is then transformed to a set of coordinate points with MDS. Researchers have advocated this method over pairwise similarity ratings based on its improved correlation with category related measures such as typicality ratings and response time (Dry & Storms, 2009; Vanpaemel, Verbeemen, Dry, Verguts, & Storms, 2010).

Using feature lists, however, can cause problems for modeling between-group differences. One concern is that dimensions that are important in certain contexts may not be generated in other contexts and therefore may be missed. Wu and Barsalou (2009) found that when participants were asked to list the properties related to a concept, participants' responses depended on how they mentally simulated that concept. Similarly, Murphy and Medin (1985) argue that concepts have an infinite set of features whose importance relative to their category membership depends on the context in which they are placed. Finally, dimensions based on feature lists incorporate categorical judgments about the values of those features, which is problematic when feature values are continuous. Some participants may not list features with continuous values, so a dimension for those features will not be created. For example, a dimension of "wings" is based on discrete feature values like absent or present, but in the case of a flying squirrel—a creature that only moderately has wings—some participants may not list wings as a feature, so a dimension of "wings" will not be created.

Modeling Rock and Janoff-Bulman's (2010) data allows for an exploration of more objective methods of creating a similarity space. Instead of subjective judgments, Latent Semantic Analysis (LSA, Deerwester, Dumais, Furnas, Landauer, & Harshman,

1990; Landauer, Foltz, & Laham, 1998) will be used to generate a similarity space and the best MDS scaling result will be identified by comparing it to Rosch's (1975) typicality data.

Method

The Original Study

The data used in Experiment 2 come from Experiment 1 of Rock and Janoff-Bulman (2010). Their study was conducted to identify differences in the effect of approach and avoidance mode on the categorization decisions of participants from across the political spectrum. First, participants were primed with either approach or avoidance mode using questions about movies or moral behavior. Participants in a neutral condition had no prime. After priming, participants were asked to make a series of categorization decisions for 60 items from five natural language categories: Carpenters' Tools, Clothes, Furniture, Vehicles and Weapons. These items were taken from a larger stimulus set used by Rosch (1975) and typicality ratings from Rosch's study were used to select four highly typical exemplars, four moderately typical exemplars, and four atypical exemplars for each of the natural language categories. While the current analyses will be performed on each natural language category separately, Rock and Janoff-Bulman summed responses across categories and used total number of "Out-of-the-category" judgments as their dependent variable. As can be seen in Figure 7, they found that for participants in avoidance mode, exclusivity increased with conservatism, but they did not find this effect for participants in approach mode.

The Model

Rock and Janoff-Bulman's (2010) data will be reanalyzed by fitting the GCM to each natural language category. The GCM determines the probability of placing a to-be-categorized item in a given category by its similarity to exemplars. Similarity between the item and members of that category increases the probability of inclusion and its similarity between the item and exemplars not in the category decreases the probability

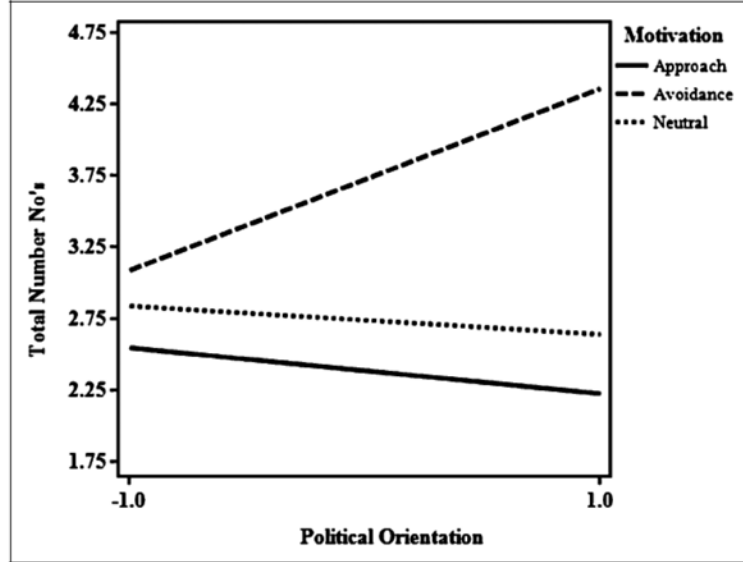


Figure 7: Figure 1 from Rock and Janoff-Bulman (2010) showing the results of their Experiment 1 for moderately typical items.

of inclusion (Equations 5 and 6). As discussed in the Introduction, a possible explanation for the effects of approach and avoidance on categorization is that being in approach mode increases the categorizer's consistency. That is, instead of probability matching, a consistent categorizer will almost always place an item in a given category if there is a greater than chance probability that it is a member. This idea was formalized in the GCM by Ashby and Maddox (1993), who modified the calculation of similarity (Equation 3) with the inclusion of γ , termed the determinism parameter,

$$S_{iA} = \left[\sum_{j \in A} \eta_{ij} \right]^{\gamma} \quad (10)$$

For a categorizer with a small γ , the model predicts that items similar to both the In-the-category exemplars and the Out-of-the-category exemplars will be categorized probabilistically. For a categorizer with a large γ , the model predicts that these items will always be categorized as either In-the-category or Out-of-the-category, depending on which is considered more likely. As an example, imagine deciding whether a yacht is a Vehicle. Assume that Yacht has a sum similarity to all vehicle exemplars equal to 2 and a similarity to non-vehicle exemplars equal to 1. When γ equals 1, the GCM predicts that a

yacht will be categorized as a vehicle $\frac{2}{2+1} = 66\%$ of the time. When γ equals 2, the GCM increases the probability to $\frac{2^2}{2^2+1^2} = 80\%$ and when γ equals 3, the GCM increases the probability to $\frac{2^3}{2^3+1^3} = 89\%$.

Populating the Model with Exemplars

At the heart of an exemplar model is the idea that a to-be-categorized item is compared to a set of exemplars. The version GCM fit in Experiment 2 will be populated with a list of exemplars that are potential category members and exemplars associated with the to-be-categorized items. This is based on the assumption that if you ask a person if an item is a Vehicle, they are at most likely to compare it with:

1. Exemplars that are known to be vehicles. Other known category members have been shown to impact unrelated categorization decisions (Brooks, Norman, & Allen, 1991);

2. Exemplars that are similar to vehicles but are not vehicles. An item's similarity to members of a contrasting category has been shown to impact categorization decisions about that item (Verheyen et al., 2011); and

3. Exemplars that are associated with the item through common usage or lexical effects. Associations between words that are lexical neighbors have been shown to impact categorization decision times (Rodd, 2004).

As mentioned above, Rock and Janoff-Bulman (2010) selected their stimuli as a subset of stimuli used in a study by Rosch (1975). Since Rosch's original category lists contained many more items—some highly typical of their category, some highly atypical, and some related but not in the category—the remainder of these lists were added to the GCM as exemplars (x_j) to flesh out the similarity comparisons of the to-be-categorized items (x_i). The shortest list (Vehicles) had 47 items while the longest lists (Carpenters' Tools and Weapons) had 60 items.

In addition to potential category members, a to-be-categorized item could bring to mind exemplars that are associated through common usage. For instance, when a

couch is Furniture, a categorizer may bring to mind a psychiatrist, even though hopefully a psychiatrist is not in danger of being mistaken for Furniture. Therefore, for each stimulus used by Rock and Janoff-Bulman (2010), up to three words frequently named during a word association task (Nelson, McEvoy, & Schreiber, 1998) were included in the exemplar list. The rules for selecting associated words were: 1) they could not be members of the category in question, 2) they were a concrete noun, and 3) they could not be related to the target by an IS-A, HAS-A, MADE-OF, or IS-A-TYPE-OF relationship. These rules assured that every item considered was on the same taxonomic level and would not lead to awkward questions, like “Is an engine a vehicle?” Such questions would not be answerable without more information, like “What object does the engine help to make up?” Only exemplars at the same taxonomic level are useful for making a categorization decision.

Approximating Category Membership

According to the GCM, a to-be-categorized item is compared to all other exemplars whose category membership is known. This requires a measure of whether an exemplar is a member of a given category. WordNet (Princeton University, 2010), an Internet based dictionary, provided such a measure. WordNet is unique in that it represents the meaning of a word both through a dictionary definition as well as a series of taxonomic relations between words as determined by lexicographers.

For the purposes of the model, exemplars were considered a member of a given category if either their dictionary definition or their taxonomy mentioned the category. As an example of the later, according to WordNet, a car is a motor vehicle, which is a self-propelled vehicle, which is a wheeled vehicle, which is a vehicle. Therefore, a car is considered a vehicle for the purposes of the GCM. It is defined, however, as, “a motor vehicle with four wheels.” This definition would also be sufficient for it to be considered a vehicle for the purposes of the GCM. This method, however, presents a problem. Two of the categories used by Rock and Janoff-Bulman (2010) did not appear reliably

indexed and therefore cannot be fit by the GCM using this method. There is no category of Carpenters' Tools in Wordnet. Additionally, weapons appear to be haphazardly labeled. For instance, "knife" has two definitions: one in the taxonomy for Tool and one in the taxonomy for Weapon, but "switchblade" appears only in the taxonomy for Tool. Fitting a prototype model as a method to address this will be discussed in the Additional Models section. Two potential members of the category Furniture are not represented in Wordnet ("end table" and "night table") and were removed from the list of exemplars.

Many of the words used as exemplars are polysemous, and category membership depends on which meaning is selected. For example, a table is not only a piece of furniture, but a format for arranging information. Therefore, definitions were selected based on the following criteria: 1) if any one of the definitions was a member of the category in question, that definition was used, and 2) if multiple definitions were members of the category in question (e.g., a jumper is a sweater, a children's cover-all, and a jacket), the first definition was used.

Determining the Similarity Space

The GCM requires all exemplars to be represented in a multidimensional similarity space. A more objective way to generate a similarity space than methods previously used to fit the GCM to natural language categories (Dry & Storms, 2009; Vanpaemel et al., 2010) is to compute measures of semantic distance based on LSA (Deerwester et al., 1990; Landauer et al., 1998). LSA uses collections of digitized texts and defines a word's meaning as the index of the various passages in which it occurs. The similarity between two words can be calculated by the overlap of the contexts in which they co-occur (as well as co-not-occur). Using a variety of passages for LSA ensures that words will occur in a variety of contexts so that context-dependent similarity can be captured. For the purposes of fitting the GCM, passages were taken from books on the website Project Gutenberg that were selected as being appropriate for college freshmen.

Again, these three measures were evaluated on their ability to predict Rosch's typicality data. Rosch (1975) herself defines an item as being typical of its category when it is similar to category members and dissimilar to non-category members and this theory has been verified empirically (Verheyen et al., 2011). An approximation of category typicality was generated for each of the three measures by summing the similarity of each exemplar to all category members (as defined by WordNet) and subtracting its similarity to non-category exemplars. This produced three typicality measures for each exemplar: one for each similarity measure.

Converting Similarity to Distance

The next step to fitting the GCM was to convert the similarity space to geometric locations using MDS. When this technique was applied in Experiment 1, the optimal number of dimensions for the solution was a priori known to be four: one for each dimension of the jellyfish. For a natural language category, however, the correct number of dimensions needs to be inferred. Therefore, a series of MDS solutions were generated with up to 20 dimensions. For each solution, an item's distance from category members (as determined from WordNet) was subtracted from that item's distance from non-category members (as determined from WordNet) to create a measure of that item's typicality. A typical item would be distant from non-category members and close to category members and score highly on this measure.

These typicality measures were correlated with Rosch's (1975) typicality data. The optimal dimensional solution was identified by locating the "elbow" in the correlation coefficients, where adding additional dimensions does not notably improve correlation. A similar criterion for fitting Prototype models to natural language categories was advocated by Verheyen, Ameel and Storms (2007). Optimal dimensionality and correlation between the generated measure of typicality and Rosch's typicality data are given in Table 3 and graphs of these correlations for all dimensions are given in Appendix Figure D2.

Table 3
Optimal Number of MDS Dimensions and Correlation with Measured Typicality by Category

	Clothing	Furniture	Vehicles
<i>r</i>	0.67	0.59	0.48
# of Dimensions	4	7	9

Generating the Data

While the GCM is designed to predict probabilistic categorization behavior, Rock and Janoff-Bulman (2010) only collected one categorization decision per item from each participant. This may be a necessity when using natural language stimuli since asking participants about the same stimulus multiple times may generate demand characteristics for participants to change their answers. Participants who are asked if an olive is a fruit many times may think that the experimenter is unhappy with their first response. Therefore, in order to generate data that can be fit by the GCM, categorization responses were combined across participants to derive the proportion of participants who placed a given item in the category.

First, participants were grouped by political leaning. Participants in this experiment were asked the same political identity questions used to identify eligible participants to Experiment 1 and political identity scores were calculated the same way. Participants had a mean political identity score of 3.47 and a standard deviation of 0.62. Liberals were identified as participants whose political identity score was one or more standard deviation below the mean. Conservatives were identified as participants whose political identity score was one or more standard deviation above the mean. Participants with missing or obviously incorrectly entered data were removed from the analysis. Out of the remaining 139 participants, 66 were identified as liberal and 73 identified as conservative. Due to the relevance of political leanings to Rock and Janoff-Bulman's conclusions, politically neutral participants were excluded from this analysis. Forty-seven participants were primed with Approach mode (Lib = 25, Con = 22), 59 with Avoidance

mode (Lib = 24, Con = 35), and 33 were in the Neutral Condition (Lib = 17, Con = 16). Participants in the moral and amoral primes were combined together, since no difference between these groups was identified in the original study. Participants' categorization decisions were combined within their groups to create the data modeled by the GCM.

Results

Since Rock and Janoff-Bulman's (2010) conclusions concerned Liberals and Conservatives in approach or avoidance condition, only these four data sets were modeled. This provided 48 data points to each model (12 to-be-categorized items \times 4 groups). The models were fit as follows. First, a fully constrained model was fit that used the same c , w , and γ parameters for each of the four data sets. This model was run 100 times with random starting points, and the parameters of the best fitting model were kept. The models were evaluated based on the Akaike Information Criterion (*AIC*, Burnham & Anderson, 2002). *AIC* uses the log likelihood of a predicted model having produced the observed data and then penalizes this measure for every parameter in the model. This method of measuring a model's fit is appropriate in Experiment 2 because it includes a likelihood measure for binary categorization responses (Wickens, 1982).

Next, a series of models that allowed between-group variability in parameters were fit. For each parameter type (c , w , and γ), three types of model were fit. The first model type allowed the given parameter to vary between participants in approach and avoidance mode. The second model type allowed the given parameter to vary between Liberal and Conservative participants. The third model type allowed the given parameter to vary between Conservatives in avoidance mode and all other participants since this was the significant interaction identified by Rock and Janoff-Bulman's (2010). Each model was fit 100 times, each time starting with the best fitting parameters from the fully constrained model permuted with random numbers sampled from the normal distribution $N(0,1)$. The fully constrained model was rerun with the same method of generating starting values as the other models. Therefore, ten models in total were tested: three

between-group contrasts for each of the three parameter types, plus the fully constrained model. Again, *AIC* was used to select the best fitting model from among these 10 models. The parameter values and *AIC* scores for all models are shown in Appendix F.

Clothing

For clothing, the best fitting fully constrained model (five parameters-3 w , c , and γ) fit with an *AIC* = 87.8. This model was improved on by allowing γ to vary between participants in the approach and avoidance conditions (six parameters), *AIC* = 86.8. The best fitting model fit both conditions with a c of 31.08, a w_{range} of .80 and fit participants in approach mode with a γ of 5.06 and participants in avoidance mode with a γ of 4.03. The parameter values and *AIC* scores for all models are shown in Appendix F. Participants in approach mode were best fit with a higher gamma than participants in avoidance mode, indicating that they are more consistent than participants in avoidance mode. The range of w weights, .8, indicates that the majority of attention weight was placed on one dimension. Figure 8 shows participant performance paired with model predictions.




Figure 8: Best fitting model predictions for Clothing alongside participants' responses

Furniture

For furniture, the best fitting fully constrained model (8 parameters-6 w , c , and γ) fit with an $AIC = 215.3$. This model was improved on by allowing γ to vary between Liberals and Conservatives (9 parameters), $AIC = 215.0$. The best fitting model fit both conditions with a c of 20.14, a w_{range} of .51 and fit Liberal participants with a γ of 4.33 and Conservative participants with a γ of 3.58. The parameter values and AIC scores for all models are shown in Appendix F. Liberals were best fit with a higher γ than Conservatives, indicating that they categorized more consistently. The w weight range of .51 implies that participants spread their attention across more dimensions for Furniture than they did for Clothing. Figure 9 shows participant performance paired with model predictions.

Vehicles

For vehicles, the best fitting fully constrained model (10 parameters-8 w , c , and γ) fit with an $AIC = 311.9$. This model was improved on by allowing γ to vary between




Figure 9: Best fitting model predictions for Furniture alongside participants' responses

Conservative participants in the avoidance mode compared with all other participants (11 parameters), $AIC = 305.0$. The best fitting model fit both conditions with a c of 8.05, a w_{range} of .52 and fit Conservative participants in avoidance mode with a γ of 9.97 and all other participants with a γ of 8.73. The parameter values and AIC scores for all models are shown in Appendix F. Conservative participants in avoidance mode were best fit with a higher gamma than Liberal participants or Conservative participants in approach mode. Again, a w weight range of .51 implies that attention was distributed across multiple dimensions.

Figure 10 shows participant performance paired with model predictions. Note that the fits for wheelbarrow and go-cart show more error than other fits so far presented. For wheelbarrow, this is because WordNet represents wheelbarrows as Vehicles, and the model is unable to totally disregard this representation when participants do not agree. While the model can represent uncertainty about a to-be-categorized item based on the exemplars that it is similar to, it would be an overly flexible model if it could account for




Figure 10: Best fitting model predictions for Vehicles alongside participants' responses

any categorization performance with a given similarity space. As for go-carts, WordNet does not represent them as Vehicles while participants believe that they are. This may be due to confusion between go-carts (strollers) and go-karts (small racing vehicles).

Additional Modeling Considerations

Alternate Models

This section presents two alternate models that were tested but not included in the final results because they did not provide a better fit (as measured by *AIC*) than the more basic model originally described. The first addresses the effect of exemplar frequency. In some versions of the GCM (e.g., Nosofsky, 1988), the similarity between a to-be-categorized item and an exemplar (Equation 3) is weighted by frequency of the exemplar, represented by the equation,

$$S_{iA} = \left[\sum_{j \in A} N_j \cdot \eta_{ij} \right]^{\gamma} \quad (11)$$

where N_j represents the number of times item j was shown during the experiment. In some category learning experiments, stimuli are presented at different rates and this is shown to moderate their impact on future categorization decisions.

Even when people are familiar with items with before they enter the experiment, they may be much more familiar with some items than others. It is unlikely that Conservatives and Liberals have substantially different kinds of encounters with objects in their daily lives (e.g., if conservatives saw more ashtrays than liberals), which could cause differences in categorization. Instead, it may be important to include an exemplar frequency weight to account for between-item differences in categorization. For example, people may directly associate chairs with Furniture, and so the chair exemplar would wholly drive the categorization decision. On the other hand, when people categorize ashtrays they may include category information from many other items, since people see few ashtrays and therefore an ashtray's own category information carries little weight.

These frequency weights do not need to be free parameters (that is, fit to the categorization data) since every exemplar needs one estimate. Free parameters would

greatly increase the complexity of the model. Instead, frequency weights can be approximated from a separate data set. While it is impossible to know for certain how often University of Massachusetts undergraduates encounter chairs and ashtrays, word frequency counts from a large corpus such as Subtlex (Brysbaert & New, 2009) are a good approximation. RT in a lexical decision task has been shown to be (relatively) sigmoidally related to $\log_{10}(\text{word frequency/million})$ (Emmanuel, Kevin, & Marc, 2010). Based on the assumption that time to make a lexical decision reflects the accessibility of an exemplar, this log transformation was used to approximate exemplar frequency (N_j in Equation 11). As examples, a low frequency exemplar such as *divan* received a weight of .28 while a high frequency exemplar such as *sofa* received an activation weight of .68. Multi-word exemplars, such as *high fidelity stereo system*, were not indexed in Subtlex, and were approximated with the average \bar{N}_j .

For the most part, including frequency information neither improved the model fit nor qualitatively changed the results. When frequency information was included for Clothing, the best fitting model also allowed γ to vary between participants in the approach and avoidance conditions, $AIC = 81.53$ (as compared to $AIC = 86.8$ without it). When included for Furniture, the best fitting model also allowed γ to vary between Liberal and Conservative participants, $AIC = 252.18$ (as compared to $AIC = 215.0$ without it). When frequency information was included for Vehicles, the best fitting model allowed c to vary between Conservatives in avoidance mode and all other participants, $AIC = 315.6$ (as compared to $AIC = 305.0$ without it).

For two out of the three categories (Furniture and Vehicles), the fit was worsened by the inclusion of frequency information. For Clothing, the improvement in fit gained by including frequency information was modest. Additionally, for Clothing and Furniture, the best fitting models allowed the same parameters to vary as without frequency information. For Vehicles, the best fitting still allowed parameters to vary between Conservatives in avoidance mode and all others, but now it identified c as the

distinguishing variable. These differences in model performance are not a significant enough improvement to justify including a model with frequency weights in the results or in future attempts to fit exemplar-based models of categorization to natural language categories.

The second alternate model considered was a prototype-based model. Since a prototype model determines category membership based on the similarity of a to-be-categorized item to a category prototype, using a prototype model avoids the requirement of knowing the category membership for any exemplars, making it possible to model Weapons and Carpenters' Tools. A prototype model proposed by Nosofsky (1987) provides a series of equations similar to the GCM. In this model, distance between the to-be-categorized item and a prototype of category A is given by,

$$d_{iPA} = \sqrt{\sum_m w_m |x_{im} - P_{Am}|^2} \quad (12)$$

where P_{Am} is the value of the prototype for category A along dimension m . Similarity between item i and prototype P_A is given by,

$$\eta_{iPA} = \exp(-c \cdot d_{iPA}) \quad (13)$$

Similarity to the prototype directly determines the probability of placing the item into category A,

$$p(A | i) = \frac{\eta_{iPA}}{\eta_{iPA} + \delta} \quad (14)$$

where δ is a free parameter representing a threshold level of similarity to a prototype that an item must achieve to be placed in that prototype's category. Note that there is no γ in this model, since it would be mathematically conflated with the c parameter.

To fit this model, the similarity between all words from Rosch's (1975) list and their category label (representing the prototype) were found using LSA, and the best number of dimensions for the MDS solution was determined by correlating Rosch's typicality results with the distance between each word and its category prototype. While this MDS solution was generated with the same exemplars used in the GCM model, only

the distances between the to-be-categorized items and the prototype were used in the model. The other words served only to provide context for the MDS solution. Parameters were allowed to vary across conditions in exactly the same way as for the GCM and all models were fit 100 times, using randomly permuted parameters from the best fitting fully constrained model as starting points. For each of the five categories, model fit information is given in Table 4.

The prototype model performed worse than the GCM by a number of measures. First, for all three categories in which there are fit measures for both the GCM and prototype model, the prototype model provides a worse fit. Second, there is much less consistency across the best fitting models. For each parameter in the prototype model, at least one category is best fit by allowing it to vary between groups. Finally, the c parameter is more difficult to interpret due to its conflation with γ in this model. This combination of factors makes the results of the prototype model less informative than the results of the exemplar model, even if its structure allows more data to be modeled.

Falsifiability

An important measure of a model is its falsifiability (Wills & Pothos, 2011). A model that can account for any result is overly general and does not do a good job of explaining any phenomena. The falsifiability of the GCM was tested by randomly reordering the data points across test items and refitting the fully constrained model to the resulting data sets. That is, the percentage of in-category responses for Cars might be randomly assigned to Elevators, while Cars may get the percentage in-category responses

Table 4
Results from Best Fitting Prototype Models

	# of Dimensions	Best Fitting Model	AIC
Carpenters' Tools	3	Conservative/Avoidance vs. Others - δ	317.8
Clothing	3	Conservative/Avoidance vs. Others - w	131.7
Furniture	3	Approach vs. Avoidance - c	257.3
Vehicles	3	Fully Constrained	452.3
Weapons	5	Conservative/Avoidance vs. Others - δ	320.4

for Camels. The model fitting proceeded exactly as described before: for each randomly ordered data set, the model was fit 100 times, each from a different randomly generated starting point. Fits were evaluated with *AIC* and the best fit of the 100 was retained. Results are reported as *AIC* values averaged across the ten random orders.

For Clothing, the average *AIC* of the best fitting models was 1316.3, for Furniture, the average *AIC* of the best fitting models was 1273.7, and for Vehicles, the average *AIC* of the best fitting models was 615.5. All these fits are substantially worse than the best fitting models when the data is correctly ordered, indicating that this model can be falsified. When 95% of participants say that a bench is a member of Furniture, but only 15% of participants say that a bed is a member of Furniture, the model cannot fit these data. It is interesting to note that the GCM fit the reordered Vehicles better than the other two categories. This is likely due to the much more graded nature of the Vehicles categorization responses as compared to Clothing or Furniture, where swapping the categorization responses would have a more substantial effect.

Discussion

The GCM provided a very good fit of category membership ratings for a certain commonly used natural language categories when supplied with category membership data from WordNet and a similarity matrix from LSA. As shown in Figure 8 through Figure 10, GCM predictions closely matched participant responses. Additionally, the Falsifiability analysis indicated that results are due to a conjunction between participant responses, objective measures of category membership and calculated measures of concept-to-concept similarity as opposed to the model's ability to fit any categorization pattern.

Other models tested did not perform better than the GCM. While a prototype model allowed for more categorization decisions to be modeled, the fits it provided were worse than those provided by the GCM. Both the quantitative measures of fit and a qualitative lack of consistency as to which parameter explains between-group

variance makes the GCM a more useful model than the prototype model tested here. Additionally, a GCM that including frequency information did not improve model fits. Perhaps participants were familiar enough with each item presented (excepting go-carts) that exemplar frequency does not impact their use in categorization decisions. While word frequency is shown to have an effect on the response time of lexical decision tasks (Emmanuel et al., 2010), participants performed the lexical decision task under tightly controlled timing conditions while participants in Rock and Janoff-Bulmman's (2010) gave untimed responses. Perhaps even a relatively uncommon exemplar will come to mind given enough time during a categorization task.

Since the dimensions relevant to natural language categorization are not known a priori and could be infinite (Murphy & Medin, 1985), it is not possible to identify the individual dimensions returned by MDS and orient the resulting coordinate points to match up with those dimensions. As in other studies in which natural language categories have been modeled (e.g. Dry & Storms, 2009; Vanpaemel et al., 2010; Verbeemen et al., 2007; Voorspoels et al., 2008) this is not a concern since the measure of interest is how attention is distributed across dimensions rather than which individual dimensions participants attended to. One way to allay any remaining concerns is to simulate categorization decisions with stimuli of known dimensionality. The GCM could then be fit to the simulated data using two different similarity spaces-one with unrotated dimensions and one with rotated dimensions. If both results show a similar spread of attention across dimensions, the current method is no cause for concern.

While this experiment has shown that mathematical models can successfully handle natural language categories, the results of these models are inconclusive. The parameters from the best fitting models of Clothing and Furniture category judgments show that Liberals and those in Approach mode are more consistent than Conservatives and those in avoidance mode. This presents an alternate interpretation of data that indicate that Conservatives in avoidance mode are more exclusive in their categorization:

if there is a greater than 50% chance that an item should be a category member, Liberals and those in approach mode are more likely to include it. The parameters from the best fitting model of Vehicle category judgments, however, indicate the opposite. Conservatives in avoidance mode are more consistent than other participants.

Approach mode has been linked to political liberalism (e.g., Janoff-Bulman et al., 2008) which may explain why Liberals are consistent categorizers of Clothing and participants in approach mode are consistent categorizers of Furniture. It is surprising, however, that Conservatives in avoidance mode appeared to be consistent categorizers of Vehicles. Note that these data were analyzed very differently-averaging across participants rather than stimuli and selecting out participants-than in Rock and Janoff-Bulman's original analysis. Additionally, recall from the Literature Review that there is dispute about whether approach mode broadens cognition relative to avoidance mode (Friedman & Förster, 2005) or if an extreme mode in either direction narrows cognition (P. A. Gable & Harmon-Jones, 2008; P. Gable & Harmon-Jones, 2010; Price & Harmon-Jones, 2010).

It is noteworthy that Vehicles showed a much more graded category membership profile than either Clothing or Furniture. This may be due to the nature of the categories, although graded membership is typically found in all categories, especially man-made artifacts (McCloskey & Glucksberg, 1978; Verheyen, Heussen, & Storms, 2011). On the other hand, it may be due to the method for selecting stimuli for the experiment. Rock and Janoff-Bulman tried to select items with a range of typicality, as judged by data collected in Rosch in 1975. It is possible that items (such as ash-trays) have significantly changed their typicality in the last 35 years, and these data are no longer useful for selecting stimuli. New typicality data for Furniture and Vehicles will be collected in a pilot study of Experiment 3.

Experiment 3 will extend these results by collecting category judgment decisions for Rosch's full lists of Furniture and Vehicles. Collecting a larger list of potential

category members will assure a full range of category typicality (that is, items that are clearly in the category, items that are clearly not in the category, and items that are in-between). It will also provide many more degrees of freedom for each model, since increasing the number of data points predicted does not increase the number of parameters required by the model. Furthermore, item presentation will be computerized, allowing for response time data to be collected and both modeled with the EBRW. The techniques shown to best work for the GCM in Experiment 2 will be directly applied to the EBRW in Experiment 3.

The techniques tested by modeling Rock and Janoff-Bulman's data proved to be successful in accounting for categorization performance. The results of these analyses, however, should not be seen as a refutation of Rock and Janoff-Bulman's results. In this analysis responses were combined across participants (as opposed to across items, as in the original analyses) and many participants were removed from the analysis to reduce the modeling computations. The main goal of these analyses was to determine the best method for fitting exemplar-based models of categorization to natural language data when looking for between-group differences so that this method can be applied to data collected in Experiment 3. Using one set of data to establish a model and method and then applying that model to a second set of data helps assure that the model is fitting the signal of the data and not the noise.

CHAPTER 6

EXPERIMENT 3

Experiment 3 is a partial replication and extension of Rock & Janoff-Bulman (2010). In Experiment 3, Liberal and Conservative participants in approach and avoidance mode make categorization decisions for Clothing and an extended list of Furniture and Vehicles while using a computerized method that allows for the collection of response times. First, this extension increases the data point to parameter ratio, providing more confidence in model results. Second, the collection of response times allows for the EBRW to be fit to these data. The EBRW will be used to account for the between-group differences. The EBRW may provide a better analog than the GCM for the differences in cognition observed between people in approach and avoidance mode.

In addition to the main study, a norming study was conducted to collect typicality and familiarity measures of the Furniture and Vehicles stimuli. Since Rosch (1975) generated these lists more than 35 years ago, some items may have become more or less typical in the intervening years or have become unfamiliar altogether. For example, consider the disappearance of television cabinets over the past decade due to the popularity of plasma televisions.

Method

Participants

Experiment 3 was run concurrently with Experiment 1 and participants were recruited through the same methods (see Experiment 1 for details). Participants self-selected to participate in Experiment 1 or in this experiment, and received one extra-credit point for participating. Over three semesters, 118 University of Massachusetts, Amherst students participated in Experiment 3 (32 Males and 86 Females). Of these 118 participants, 64 were Liberals and 54 were Conservatives. 22 Liberals and 17 Conservatives were in the Approach condition, 21 Liberals and 19 Conservatives were

in the Neutral condition, and 20 Liberals and 19 Conservatives were in the Avoidance condition. One participant was not included in analyses due to computer error.

The norming study examining familiarity and typicality of the stimuli had 49 participants, drawn from the psychology department subject pool in the Fall of 2011. There was no political orientation requirement to participate in the norming study. Furthermore, none of the participants in who participated in the norming study also participated in Experiment 1 or Experiment 3. Of these participants, 39 were Females and 10 were Males.

Materials

As previously mentioned, all categorization stimuli originated from Rosch (1975). The category of Clothing was used as a practice set and the twelve-item subset tested by Rock & Janoff-Bulman (2010) were used in the present study. For Vehicles and Furniture, all items from Rosch's list were used. All stimuli were presented on an eMac running the Psychophysics Toolbox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997). For the norming study, only the Vehicles and Furniture lists were tested.

Procedure

Experiment 3 had three phases: approach/avoidance induction, categorization, and manipulation check. After reading and signing the consent form, participants sat at a computer in an isolated room and performed the approach/avoidance induction as described in Experiment 1. Participants in the approach condition were asked to list 10 movies to watch for a good time, participants in the avoidance condition were asked to list 10 movies to not watch to avoid a bad time, and participants in the neutral condition were asked to list 10 movies they had watched recently. All participants were given a maximum of 5 minutes to complete this stage.

After listing movies, participants were asked to make a series of categorization decisions. Each participant started with twelve Clothing categorization decisions in order to familiarize them with the response format (since response time is critical to the

EBRW). On each trial, participants were shown an item (e.g., “shirt,” “sweater,” “ring”) and asked whether it was a member of the category “Clothing.” If participants believed that a given item was in the category, they pressed the key “f,” otherwise they pressed the key “j.” The presentation order of the twelve Clothing items was randomized for each participant. After categorizing Clothing, participants were asked to make categorization decisions about Vehicles and Furniture. Both the order of category presentation and the order of item presentation were randomized across participants.

After making categorization decisions for all three sets of stimuli, participants performed the manipulation check as described in Experiment 1. They were asked to list three goals that they hoped to accomplish and three goals that they felt obligated to accomplish. Response times to the former are believed to reflect the extent of approach mode activation while response times to the later are believed to reflect the extent of avoidance mode activation. The order of presentation of the two goals sets was randomized across participants.

Participants in the norming study were shown the same lists of Furniture and Vehicles as participants in Experiment 3. Instead of making categorization decisions for each item, however, participants were asked to judge their familiarity with each item as well as the item’s typicality as a member of its category. Each judgment was made on a scale of 1 to 9, where 1 meant Not Familiar or Highly Atypical and 9 meant Highly Familiar or Highly Typical. Participants were encouraged to use the whole range of responses.

Results

Norming Study

The norming study aimed to answer two questions. First, how familiar were participants with the natural language category members used as stimuli in this study? Average familiarity ratings are listed in Appendix Table F8 and Table F9. Overall, participants were very familiar with the stimuli in the study. On a scale of 1 (Not

Familiar) to 9 (Very Familiar), the Furniture stimuli had a median familiarity rating of 8.79, $IQR = 0.76$ and the Vehicles stimuli had a median familiarity of 8.55, $IQR = 0.72$. Items rated as unfamiliar were excluded from further analyses. Since both familiarity distributions were highly negatively skewed, outliers were identified with the $1.5 \times IQR$ method. For Furniture, this led to the exclusion of Davenport, Divan, Cedar Chest, Chaise Lounge, Hassock and Hi-Fi. This method did not lead to the exclusion of any Vehicles

The second question addressed by the norming study was how contemporary participants ranked the typicality of these stimuli as members of their natural language categories. Average typicality ratings are listed in Appendix Table F8 and Table F9. While Rosch (1975) had collected this data thirty-seven years ago, it was possible that the typicality of some of these items had changed as the culture has evolved. Overall, correlation was very high for Vehicles ($r = -0.93$) and high for Furniture ($r = -0.72$). When unfamiliar items are removed from Furniture, correlation increased to $r = -0.86$.

Manipulation Check

As in Experiment 1, the measure of interest for the manipulation check was the amount of time participants required before coming up with each of the three goals they aspired to accomplish and each of the three goals they were required to accomplish (Friedman & Förster, 2001). This was quantified by the average amount of time participants required before they started typing each of their answers. Differences in the time to start typing were tested by a $3(\text{Mode}) \times 2(\text{Question Type})$ mixed model ANOVA, with Question Type as a within-subjects factor. First, there was a main effect of Question Type. Participants were faster to start typing goals that they aspired to achieve ($\bar{x} = 6.12, s = 4.24$) than goals that they were required to achieve ($\bar{x} = 10.39, s = 8.36$), $F(1,114) = 33.70, p < .001$. There was, however, neither a main effect of Mode ($F(1,114) = 0.35, p = .70$) nor an interaction between Mode and Question Type ($F(2,116) = 0.64, p = .53$).

Categorization Performance

The goal of the present analyses is to identify between-group differences in categorization decisions as reported by Rock and Janoff-Bulman (2010). One participant was removed from further analyses for answering “in-the-category” to 56 of 58 Furniture, 45 of 47 Vehicles and averaging 0.23 seconds per answer for Furniture. Average “in-the-category” rates are given in Table 5. Complete “in-the-category” rates are provided for each item in the Appendix Table F8 and Table F9.

Table 5
Average “In-the-category” Rates for Experiment 3

Condition	\bar{x}	<i>s</i>
Liberals		
Approach	.60	.11
Neutral	.56	.09
Avoidance	.61	.09
Conservatives		
Approach	.54	.08
Neutral	.58	.14
Avoidance	.65	.08

In Experiment 3, participants were selected for their political identity and this variable must therefore be treated as categorical. The inclusion of many category items, however, means that typicality can now be treated as a continuous variable. Categorization decisions were analyzed through a logistic regression, with Subject, Item Typicality (as measured by the norming study), Mode (Approach, Neutral, Avoidance), Political Identity (Liberal vs. Conservative), and their interactions entered as predictors of category inclusion. The addition of these predictors significantly improved the predictions of the model over a constant-only model from 59.3% correct to 77.7% correct, $\chi^2(12) = 4884.24, p < .001$. The results of this test are given in Table 6.

Table 6
Logistic Regression for Experiment 3

Source	<i>B</i>	<i>SE</i>	<i>Wald</i>	<i>df</i>	<i>p</i>	<i>odds</i>
1. Constant	-3.37	0.18	345.06	1	< .001	0.03
2. Typicality	0.76	0.04	475.35	1	< .001	2.13
3. Subjects	0.00	0.01	0.01	1	.92	1.00
4. Mode 1	-0.93	0.27	12.13	1	< .001	0.40
5. Mode 2	-1.02	0.27	14.73	1	< .001	0.36
6. Politics	-0.66	0.27	5.99	1	< .05	0.52
7. Mode 1 × Typicality	0.13	0.05	5.92	1	< .05	1.13
8. Mode 2 × Typicality	0.14	0.05	7.54	1	< .01	1.15
9. Politics × Typicality	0.04	0.05	0.60	1	.44	1.04
10. Mode 1 × Politics	1.15	0.39	8.77	1	< .01	3.17
11. Mode 2 × Politics	1.34	0.38	12.14	1	< .001	3.81
12. Mode 1 × Politics × Typicality	-0.08	0.08	1.11	1	.29	0.92
13. Mode 2 × Politics × Typicality	-0.16	0.07	4.96	1	< .05	0.85

Note: Mode 1 is Approach = 0, Avoidance = 1, Mode 2 is Approach = 0, Neutral = 1. Politics is Liberal = 0, Conservative = 1.

In this analysis, all groups were compared to Liberals in approach mode. The categorization rates predicted by the logistic regression are illustrated in Figure 11. For Liberals in approach mode, each unit increase in Typicality increased the odds of including an item in a given category (Source 2). For example, the regression equation predicts that Liberals in approach mode will make an “in-the-category” judgment for benches (typicality = 7.32) 89.6% of the time, pianos (typicality = 4.83) 56.8% of the time, and sewing machines (typicality = 2.60) 19.6% of the time.

For Liberals in avoidance mode, there was an overall decrease in the odds of including an item in a given category relative to Liberals in approach mode (Source 4), but this difference decreases as Typicality increases (Source 7). According to the regression equation, Liberals in avoidance mode would make an “in-the-category” judgment to benches 89.6% of the time, pianos 48.9% of the time and sewing machines 11.8% of the time. Note that for items of high typicality, the response probability is the same but for items of lower typicality the probability is lower.




Figure 11: Experiment 3 categorization rates by typicality as predicted by the logistic regression

For Conservatives in approach mode, there was also an overall decrease in the odds of including an item in a given category relative to Liberals in approach mode (Source 6) and this decrease remains for all levels of Typicality (Source 9). According to the regression equation, Conservatives in approach mode would make an “in-the-category” judgment to benches 85.7% of the time, pianos 45.2% of the time and sewing machines 12.3% of the time. Notice how the difference between Liberals in approach mode and Conservatives in approach mode continues through the whole range of typicality.

Finally, for Conservatives in avoidance mode, there was a smaller decrease in the odds of including an item in the category relative to Liberals in approach mode (Sources 4, 6 and 10) than was found for Conservatives in approach mode (Source 6) or Liberals in avoidance mode (Source 4). The difference in inclusion between Conservatives in

avoidance mode and Liberals in approach mode decreases and eventually the effect reverses as typicality increases (Sources 7 and 12). According to the regression equation, Conservatives in avoidance mode would make an “in-the-category” judgment to benches 91.3% of the time, pianos 56.4% of the time and sewing machines 16.5% of the time. Notice how Conservatives in avoidance mode are predicted to have a higher percentage of “in-the-category” responses than Liberals in approach mode and this difference only reverses at the lower levels of typicality.

Modeling

Participants’ categorization responses and response times were averaged within Mode, Political Identity and category type (Furniture or Vehicles). Each was fit with the EBRW. As with the GCM in Experiment 2, this required establishing a similarity space, approximating category membership, and determining a set of exemplars to populate the space.

Populating the Model with Exemplars

Unlike Experiment 2, participants in Experiment 3 made categorization decisions for every natural language category member on Rosch’s (1975) Vehicles and Furniture list. Therefore, no additional category members were added to the model to flesh out the exemplar list. Similar to Experiment 2, however, up to three words associated with each natural language category member were included as exemplars. Again, these associated words were taken from the University of South Florida Word Association database (Nelson et al., 1998). The full list of associated words can be seen in Appendix E.

Approximating Category Membership

As in Experiment 2, category membership was determined with WordNet (Princeton University, 2010). Exemplars whose definition or taxonomy contained the natural language category were considered to be category members for modeling purposes.

Determining Similarity Space

Experiment 3 used the same methods for determining a similarity space as Experiment 2. First, overall similarity was calculated by inputting all exemplars into LSA. The resulting similarity matrix was converted to distances with MDS. Nineteen MDS solutions were constructed, having between 2 and 20 dimensions. The optimal number of dimensions in the MDS solution was determined by correlating a constructed typicality rating for each solution and the typicality ratings collected by Rosch (1975). The optimal number of dimensions was identified by the “elbow” in the correlations, where adding more dimensions did not significantly increase the correlation. See Appendix Figure D2 for graphs of the MDS results. For both Furniture and Vehicles, the optimal MDS solution had 10 dimensions. It is likely that these differ from the optimal number of dimensions in Experiment 2 due to the increased number of exemplars used in Experiment 3.

Generating the Data

Similarly to Experiment 2, data was averaged across participants within a Political Identity, mode condition and category in order to generate proportional categorization data. As there were no hypotheses about participants in the Neutral Mode condition, these participants were excluded from the modeling analyses. Therefore, four data sets were generated for each category: Liberals in approach mode, Liberals in avoidance mode, Conservatives in approach mode, and Conservatives in avoidance mode.

Fitting the EBRW

For each category, a fully constrained version of the EBRW was fit to all four data sets. This model found the c , the set of w parameters, the A and B boundaries, and the time parameters (α , k and μ) that best minimized the $WSSD$. This model was run 100 times with random starting points to guard against local minima, and the parameter from the best fitting model were kept. Next, this model was run again 100 times with the

best fitting parameters as its starting point, each permuted by normally distributed noise ($N(0,1)$).

The next step in model fitting was to allow parameters to vary between groups. The data were fit with models that allowed either c , all w weights, A , B or both A and B to vary between groups. As in Experiment 2, parameters were allowed to vary one of three ways: between Liberals and Conservatives, participants in Approach mode and Avoidance mode, and Conservatives in Avoidance mode compared to all others. Because the logistic regression analysis of the categorization performance showed that Liberals in approach mode were the most inclusive categorizers, additional models were fit that allowed these parameters to vary between Liberals in approach mode and all other participants. Therefore, a total of 21 models were fit to account for between-group differences. Each was fit 100 times using the best fitting parameters of the fully constrained model plus normally distributed noise as starting points. The predictions of the best fitting models can be found in the Appendix Table F10 and Table F11.

In Experiment 2, models were compared using *AIC* (Burnham & Anderson, 2002), a theoretically motivated method of accounting for differences in model complexity using log likelihood as a fit measure. Since the EBRW has no known likelihood fit measure (Nosofsky & Stanton, 2005), models in Experiment 3 were compared with cross-validation (Browne, 2000). Cross-validation fits the model to one set of data and tests it on another. It is based on the assumption that a model that is too complex will fit noise in the data and do a poor job of predicting new data. Cross-validation was implemented by cycling through each to-be-categorized item, fitting each of the models with that item's categorization and RT data withheld, and then using the resulting parameters to predict the withheld data. Each model's cross-validation fit measure is the average of its *WSSD* across all predicted data points.

The best fitting parameters for each of these models are given in the Appendix Table F10 and Table F11. For Furniture, the best fitting model allowed A to vary between

Conservatives in avoidance mode and all other participants. For participant categorization decisions, and the best fitting model predictions, see Figure 12. For participant RT and the best fitting model predictions, see Figure 13. In this best fitting model, A_{ConAvo} was 0.94 while A_{Other} was 1.04, indicating that Conservatives in avoidance mode required less information to make an “in-the-category” decision and were less consistent.

For Vehicles, the best fitting model allowed both A and B to vary between Liberals and Conservatives. For participant categorization decisions and the best fitting model predictions, see Figure 14. For participants RT and the best fitting model predictions, see Figure 15. In this best fitting model, A_{Lib} was 21.44 and A_{Con} was 23.09, while B_{Lib} was 32.05 and B_{Con} was 29.51. Here, Liberals required less information than Conservatives in order to make an “in-the-category” decision and more information than Conservatives in order to make an “out-of-the-category” decision. Additionally, this places the A and B farther apart for Liberals than Conservatives, indicating that there was more consistency between Liberals than Conservatives in their categorization decisions.

Discussion

Experiment 3 showed between-group differences between Liberals and Conservatives in approach and avoidance modes when categorizing Furniture and Vehicles. When grouped across category, Liberals in approach mode showed higher levels of category inclusion relative to Liberals in avoidance mode and Conservatives in approach mode. Conservatives in avoidance mode, however, were similar in their inclusivity to Liberals in approach mode. These differences were most notable at lower levels of category typicality. When each category was modeled separately, the best fitting model of Furniture allowed Conservatives in avoidance mode to have a smaller A parameter than other participants and the best fitting model of Vehicles allowed Liberals to have a smaller A parameter and a larger B parameter than Conservatives. While the manipulation check did not find significant differences between participants in approach




Figure 12: Furniture “in-the-category” judgments and modeling predictions for Conservatives in avoidance mode and all other participants.




Figure 12 continued




Figure 12 continued




Figure 13: Furniture RT and modeling predictions for Conservatives in avoidance mode and all other participants.




Figure 13 continued




Figure 13 continued




Figure 14: Vehicle "in-the-category" judgments and modeling predictions for Liberals and Conservatives




Figure 14 continued




Figure 14 continued




Figure 15: Vehicle RT and modeling predictions for Liberals and Conservatives




Figure 15 continued




Figure 15 continued

and avoidance modes, between-group differences in inferential statistics and model results indicate that the manipulation was effective.

To the extent that cognitive differences between people with different political identities are due to personality differences, with Liberals more inclined to approach state and Conservatives more inclined to avoidance state (Janoff-Bulman et al., 2008), these results are most in line with the findings of Markman et al. (Markman et al., 2005). They found that decision boundaries are impacted by approach and avoidance state when learning a new category. For both categories in Experiment 3, between-group differences were best accounted for by changes in the decision criteria—how much information was required before a categorization decision was made.

Even though both categories were best fit with between-group differences in decision criteria, the group with the most inclusive criterion depended on the category being modeled. That is, Furniture was best fit by a model allowing Conservatives in avoidance mode had a smaller A than other participants, while Vehicles were best fit by a model allowing Liberals to have a smaller A and a larger B than Conservatives. This surprising result can be addressed by looking at the parameters of non-winning models. For Furniture, look at the model that allowed the A and B parameters to vary between Liberals in approach mode and all other participants. Liberals in approach mode were best fit with a smaller A and a larger B , the same as their categorization decisions for Vehicles. Now look at the Vehicles model that allowed the A and B parameters to vary between Conservatives in avoidance mode and all other participants. Here, Conservatives had a smaller A and an identical B , the same as their categorization decisions for Furniture. This indicates that when Conservatives in avoidance mode are more inclusive categorizers, they are also less consistent categorizers.

In sum, it seems like for both Furniture and Vehicles, Liberals in approach mode and Conservatives in avoidance mode were the most inclusive categorizers. Since Liberals are more likely to be in approach mode and Conservatives are more likely to be

in Avoidance mode (Janoff-Bulman et al., 2008), perhaps participants in both of these conditions were more receptive to their manipulation and were placed in a more extreme mode. Price and Harmon-Jones (Price & Harmon-Jones, 2010) refer to this as “high motivational intensity” and found related effects with approach motivation: that within a given mode, “low motivational intensity” and “high motivational intensity” have different effects on cognitive abilities. In contrast to the current results, however, they found that high motivational intensity led to more exclusive categorization.

While the models overall did a good job of predicting participants’ categorization decisions, some predictions were notably incorrect. Often this was due to a basic disagreement between participants and WordNet (Princeton University, 2010) about category membership. For example, while participants claimed to be familiar with a buffet, obviously few were familiar with it as a piece of furniture and the model could not accommodate giving a piece of furniture such a low probability of being in the category. Additionally, the University of South Florida Free Association Norms (Nelson et al., 1998) did not provide any associations for “buffet,” which would likely have been food and have brought down the models association between Buffets and other exemplars of Furniture.

As can be judged by the fit values in Appendix Table F10, the models of Vehicles fit worse than the models of Furniture. The two items with the worst fit were Ambulance and Skateboard. Ambulance was an odd instance. Almost all participants thought an Ambulance was a Vehicle and WordNet agreed. Ambulance’s two closest associations (other than self-similarity) were wheelchair and truck stop. These were both not considered to be Vehicles by WordNet and these associations reduced the predicted category inclusion rate to 73% for Liberals and 83% for Conservatives. Note that while these values may not have the largest absolute error, they have the largest weighted error because they have little variability according to the binomial distribution with an observed category inclusion rate close to 100%.

Skateboard had problems similar to Buffet in that there was a disagreement between WordNet and participants. Inclusion rates ranged from 15% for Liberals in avoidance mode to 38% for Liberals in approach mode. WordNet, however, indicates that a skateboard is a vehicle. After its self-similarity, Skateboard was about equally similar to Movie and Bike. Movies are not considered Vehicles by WordNet and Bikes are considered Vehicles, so these exemplars likely canceled out each other in terms of category inclusion. The model therefore predicts that Skateboards will be included in the category 83% of the time for Liberals and 80% of the time for Conservatives.

It is unclear what the best step is to improve model fits. In a number of instances, a fundamental disagreement between participants and WordNet regarding category membership causes the model to make poor predictions. While an item's WordNet category membership is obviously not the only component of the categorization decision, it is a large factor because an item's self-similarity always has the most influence on the model's prediction. One way to reduce the impact of WordNet's category membership for the to-be-categorized item is to manually reduce self-similarity. Highly typical category members will remain similar to other category members, keeping their predicted category inclusion rates high. Atypical members would not be similar to other category members and this would drive down their category inclusion rates. This would help the fits for items like Buffet and Skateboard, but not for items like Ambulance.

CHAPTER 7

GENERAL DISCUSSION

The purpose of the dissertation was to use exemplar-based models of categorization to explain the between-group differences observed by Rock and Janoff-Bulman's (2010). Experiment 1 the EBRW was fit to the categorization decisions of Conservative and Liberal participants in approach and avoidance modes using experimenter-designed stimuli. The stimuli in Experiment 1 had clearly defined dimensions that could be easily accounted for by mathematical models of categorization. In Experiment 2 the GCM was fit to the data from Rock and Janoff-Bulman's (2010) Experiment 1. In Experiment 3, the EBRW was fit to data that replicated Rock and Janoff-Bulman's (2010) for two categories and extended it by collecting RTs and adding stimuli to each of the two categories.

The Effect of Mode and Political Identity

Rock and Janoff-Bulman (2010) showed that when participants were put into avoidance mode, the most politically conservative participants were the most exclusive categorizers of natural language stimuli. Approach mode had no significant effect on the inclusivity of participants who made categorization decisions about natural language stimuli. Based on these results, one would expect that Conservatives in avoidance mode would be the most exclusive categorizers of artificial stimuli in Experiment 1 and of natural language categories in Experiment 3.

In Experiment 1, Liberals and Conservatives were placed into approach or avoidance mode while deciding whether artificial stimuli were members of a newly learned category. Experiment 1 showed that Liberals were more inclusive categorizers than Conservatives when making decisions about these artificial stimuli. Specifically, an ANOVA showed more inclusive categorization for Liberals in approach mode than Conservatives in approach mode, while a Signal Detection analysis showed that Liberals in avoidance mode had a more inclusive criterion than Conservatives in avoidance mode.

In Experiment 3, Liberals and Conservatives were placed into approach or avoidance mode while deciding whether natural language stimuli were members of natural language categories. Experiment 3 showed that Liberals in approach mode were the most inclusive categorizers of these natural language stimuli, followed closely by Conservatives in avoidance mode. At high levels of category typicality Conservatives in avoidance mode became more inclusive categorizers of these natural language stimuli than Liberals in approach mode. These results are surprising. Based on Rock and Janoff-Bulman (2010), one would expect that Conservatives in avoidance mode would be the most exclusive categorizers in both studies.

In sum, the results of the behavioral data are inconsistent across experiments. Overall, Liberals appear more inclusive categorizers than Conservatives in Experiment 1, while both Liberals in approach mode and Conservatives in avoidance mode were highly inclusive categorizers in Experiment 3. As discussed throughout this dissertation, inclusive categorization can come about through differences in a number of components of the categorization process. Differences in these components could produce different effects depending on the stimuli. If the observed between-group differences were due to differences in consistency of categorization, then the apparent inclusivity of their categorization would be dependent on the typicality of the items being categorized. In Experiment 1, the majority of the jellyfish (8 of 13) were clearly category members. Under these circumstances a more consistent categorizer would appear to be a more inclusive categorizer – more consistently including all of the typical category members. In Experiment 3, there was a greater range of typicality within each category and under these conditions being a consistent categorizer might not appear to be an inclusive categorizer. On the other hand, if being a Liberal or being in approach mode broadened a categorizer's spread of attention across dimensions, this may make them more likely to include atypical jellyfish since these jellyfish are unlikely to be atypical across all dimensions due to their creation process. Again, the effect of focused attention is harder

to predict for natural language categories. Perhaps focusing on dimension such as “does it carry people” for Vehicles or “can you buy it in a Furniture store” for Furniture may make people with focused attention appear to be more inclusive categorizers.

To address these issues, Experiment 2 used the GCM to model data from Rock and Janoff-Bulman’s (2010) Experiment 1, in which Liberals and Conservatives were placed into approach or avoidance mode while deciding whether natural language stimuli were members of natural language categories. Each of the three categories modeled indicated that differences in inclusiveness between groups were due to differences in γ , a parameter representing the consistency of responses to items of uncertain category membership. Participants in approach mode had a higher γ than participants in avoidance mode when categorizing Clothing, Liberal participants had a higher γ than Conservatives when categorizing Furniture, and Conservative participants in avoidance mode had a higher γ than all other participants when categorizing Vehicles.

Experiments 1 and 3 used the EBRW, which does not have a γ parameter, to model the data. The distance between the A and B parameters in the EBRW would be roughly analogous to γ in the GCM. Increasing the distance between A and B effectively increases the number of exemplars sampled before a decision is made and decreases the chance that the decision will be impacted by an improbable sampling of exemplars from the less-likely category. Therefore, a large distance between A and B means that the model predicts a consistent categorization decision of the most likely category membership. The absolute values of A and B determine the model’s inclusiveness. As A moves closer or B moves farther away from the zero-knowledge point, the model predicts more inclusive categorization.

Experiment 1 showed that Conservatives in avoidance mode had a larger c parameter than all other participants. This indicates that Conservatives in avoidance mode required items had to be a closer match to trained category members before they judged them to be category members. Experiment 3 showed that Conservatives in avoidance

mode had a smaller A parameter than all other participants when categorizing Furniture. This indicates that Conservatives in avoidance mode were more inclusive but less consistent categorizers than all other participants. Experiment 3 also showed that Liberals had a smaller A parameter and a larger B parameter than Conservatives when categorizing Vehicles. This indicates that Liberals were more inclusive than Conservatives. Liberals were also the more consistent categorizers, since the distance between A and B is greater for Liberals than Conservatives, indicating that Liberals required more information than Conservatives in order to reach a decision.

The exemplar-based models of categorization fit in Experiments 2 and 3 identified between-group differences as originating in the A and B parameters or the γ parameter. These parameters are all associated with the decision-making components of the categorization process. That is, they determine how people turn an item's similarity to a set of exemplars into a categorization decision, rather than how they scale the similarity between exemplars. These results are consistent with those of Markman et al. (2005), who found that being in different modes during a category learning task affected how participants established their decision boundaries. Modeling results for Experiment 1, however, identified between-group differences as originating in the c parameter, which is associated with the item-to-exemplar comparison component of categorization. It is possible that the nature of the stimuli interact with the observed effect of the categorization process. That is, approach and avoidance mode could impact the perception of items learned in the laboratory, while at the same time impacting decision making about previously known items.

Previous attempts to account for the effects of context on categorization have focused on how people spread attention across dimensions (e.g., Lin & Murphy, 1997; Spalding & Murphy, 1996). The modeling results presented here indicate that researchers should extend their investigations beyond how attention is distributed across dimensions to look at other aspects of the categorization process, such as the amount of information

required to make a categorization decision and the impact of context on the consistency of categorization responses. For example, Lamberts (1995; 1998) has shown that making categorization decisions under a time-constraint affects how information accumulates. There may be between-group differences in this effect. Other types of ephemeral and persistent contexts may also affect decision criteria.

Different analyses—ANOVA, Signal Detection, and modeling—of data from Experiments 1 and 3 yielded different results about whether the manipulations were successful. For reasons explained in the Discussion of each experiment, these different results indicate that the manipulation check used was not sufficiently sensitive to differences between participants in approach mode and participants in avoidance mode. Shortly after the current experiments were designed, a line-bisection task was proposed as an alternate manipulation check and shown to correlate with neuropsychological evidence associated with differences in mode (Nash, McGregor, & Inzlicht, 2010). Future experiments may benefit from including multiple manipulation checks that have been shown to be sensitive to approach/avoidance mode manipulation, such as line-bisection.

Modeling Natural Language Categories

Previous attempts to model natural language categories have used subjective measures of generating a similarity space and have not investigated between-group differences (e.g., Smits et al., 2002; Voorspoels et al., 2008). The present methods were designed to develop a more objective way of fitting exemplar-based models of categorization to natural language categories than has previously been used, allowing them to account for between-group differences in categorization decisions. Three major challenges needed to be overcome: defining a similarity space, determining category membership for exemplars, and populating the similarity space. In Experiments 2 and 3, LSA (Landauer et al., 1998) was used to define the similarity space of the natural language categories, WordNet (Princeton University, 2010) was used to define category membership of exemplars, and the University of South Florida Word Association

database (Nelson et al., 1998) was used to populate the similarity space of the natural language categories with exemplars. These methods provided a successful framework for modeling between-group differences in natural language categorization.

Future attempts to fit exemplar-based models to natural language categories could use other methods. For example, similarity metrics based on WordNet taxonomies could be used to define the similarity space of natural language categories rather than LSA (for a discussion of these metrics, see Budanitsky & Hirst, 2006). Unlike LSA, in which words are defined by the context in which they are found, WordNet uses synonyms to define words. WordNet therefore distinguishes between the various meanings of polysemous words where LSA does not. For example, in WordNet the Buffet of dining room furniture is unrelated to the Buffet of self-service restaurants, so concepts similar to Buffet_{Restaurant} will not affect the similarity calculations of Buffet_{Furniture}. This may have a positive or a negative impact for modeling categorization decisions, depending on whether homonyms and associated words impact the categorization decision.

Items from the University of South Florida Word Association database (Nelson et al., 1998) were included as exemplars in the similarity space of a to-be-categorized item only if those items were of the same taxonomic level as the to-be-categorized item in Experiments 2 and 3. For example, drawers and wood were not included as exemplars in the similarity space when modeling participants' decisions about whether a desk is Furniture even though these two items are associated with desks in the database. Model fits might be improved by including items as exemplars regardless of their taxonomic level. Sometimes items from different taxonomic levels that are associated with an item appear in an exemplar-based model as features that the item has in common with exemplars of a category. For example, a desk can be similar to exemplars of Furniture, like chairs and tables, because they are all made of wood. If these features were exemplars themselves, their category membership would need to be determined. The WordNet definition of wood does not mention Furniture, so the exemplar wood would

not be included in the category Furniture if WordNet is used to define the category membership of exemplars. The EBRW would then lower the probability that a desk would be included in the category Furniture based on its similarity to wood. On the other hand, the WordNet definition of drawers does mention Furniture, so the exemplar drawers would be included in the category Furniture. The EBRW would then raise the probability that a desk would be included in the category Furniture based on its similarity to drawers.

Future efforts to model natural language categorization decisions could benefit from trying to estimate parameter values from other measures. For example, bias could be estimated by a signal detection task unrelated to the categorization decision. This estimate could stand in for the distance between the A and B boundaries, so only one parameter would have to be fit to the data. Similarly, the score from a recognition memory test could be used to estimate the c parameter. Estimating these parameters from other sources would have two benefits. First, it would increase the meaningfulness of the model by showing that the parameters reflect measureable psychological processes. Second, from a modeling standpoint, it would decrease the complexity of the model and reduce the likelihood that the model is fitting noise in the data rather than true underlying behavior.

Conclusion

In conclusion, the present studies identified differences in the categorization decisions of Liberals and Conservatives in approach and avoidance modes, but these differences were inconsistent across studies and categories. These categorization decisions were fit with exemplar-based models of categorization. To do so required developing objective methods of generating model components. Liberals and those in approach mode may have a more inclusive category criterion but are most likely more consistent categorizers compared to Conservatives and those in avoidance mode. The effect of differences in consistency on category inclusiveness, however, depend on the category structure being tested.

APPENDIX A

PRE-SCREENING SURVEY QUESTIONS

1. Where would you place yourself politically on the following two scales?
 - a. 1 (very liberal) – 4 (neither) – 7 (very conservative)
 - b. 1 (strong Democrat) – 4 (neither) – 7 (strong Republican)
2. How much do you tend to like or dislike political conservatives?
 - a. 1(dislike extremely) – 7 (like extremely)
3. How much do you tend to like or dislike political liberals?
 - a. 1(dislike extremely) – 7 (like extremely)

APPENDIX B

INSTRUCTIONS FOR APPROACH/AVOIDANCE MANIPULATION IN EXPERIMENTS 1 AND 3

1. Approach: We are interested in finding movies that are generally considered fun to watch. Please list 10 movies you would recommend a friend TO SEE if they wanted to have a good time.

2. Neutral: We are interested in finding movies that are generally well known. Please list 10 movies that you have seen recently.

3. Avoidance: We are interested in finding movies that are generally considered unpleasant to watch. Please list 10 movies you would recommend a friend NOT TO SEE if they wanted prevent a bad time.

APPENDIX C

STIMULI FOR EXPERIMENT 1

Table C1.

Stimuli Generation Values for Experiment 1

Dimensions	Stimulus Space		Category Prototype Values	
	Min	Max	Mean	SD
Organ color (rgb = [x 0 0])	85	205	145	7.5
Bell diameter (400x by 200x pixels)	.6	1.5	1.05	.05
Bell thickness (line = x pt weight)	1	11.8	6.45	.6
Tentacle length (segment = x pixels)	16	34	25	1

Table C2.

Stimuli Values for Experiment 1

	Organ Color	Bell Diameter	Bell Thickness	Tentacle Length
High 1	0.99	24.90	6.68	133.92
High 2	1.03	26.43	7.10	149.54
High 3	1.16	25.29	5.85	145.24
Medium 1	1.07	23.44	6.72	149.82
Medium 2	1.11	24.51	6.29	143.42
Medium 3	1.05	25.76	5.49	138.99
Atypical 1	1.10	26.68	3.34	149.61
Atypical 2	1.00	23.89	7.17	123.07
Atypical 3	1.13	23.83	7.04	142.03
Non-category 1	0.90	30.59	2.32	204.88
Non-category 2	0.89	16.94	7.85	153.59
Non-category 3	1.14	22.51	2.65	123.53
Non-category 4	0.65	25.23	6.51	94.37
Non-category 5	0.75	29.26	6.80	159.42
Non-category 6	1.48	16.58	6.30	132.58
Non-category 7	1.26	25.79	6.17	88.19

Note. Stimuli in **boldface** were used during training

APPENDIX D
DIMENSIONAL ANALYSIS FOR MULTIDIMENSIONAL SCALING




Figure D1: Dimensional analysis for Experiment 2





Figure D2: Dimensional analysis for Experiment 3

APPENDIX E

WORD ASSOCIATIONS FOR NATURAL LANGUAGE CATEGORY MODELING

Experiment 2

Carpenter's Tools	Clothing	Furniture
1. Drill	1. Shirt	1. Drapes
a. Hole	a. Shorts	a. Window
b. Teeth	2. Cane	b. House
2. Ladder	a. Crutch	c. Wall
a. Roof	b. Walker	2. Desk
b. Stair	3. Purse	a. Pen
c. Paint	a. Money	b. Clerk
3. Bolt	b. Wallet	c. Paper
a. Door	c. Girl	3. Ashtray
4. Saw	4. Sweater	a. Cigarette
a. Tree	a. Winter	b. Smoke
5. Rag	5. Ring	c. Car
a. Towel	a. Marriage	4. Lamp
b. Dish	6. Pants	a. Light
c. Bum	a. Man	
6. Scissors	7. Bracelet	
a. Paper	a. Wrist	
7. Screwdriver	8. Pajamas	
a. Vodka	a. Bed	
b. Orange Juice		
8. Blueprint		
a. House		
9. Hinge		
a. Gate		

Vehicles

1. Car
 - a. Money
 - b. Race
 - c. Road
2. Jet
 - a. Airline
3. Camel
 - a. Desert
 - b. Water
4. Yacht
 - a. Ocean
 - b. Sea
5. Train
 - a. Track
 - b. Station
6. Tractor
 - a. Farm
 - b. Grass
 - c. Dirt
7. Bus
 - a. People
 - b. City
 - c. Station
8. Elevator
 - a. Stair
 - b. Building
 - c. Box
9. Airplane
 - a. Airport
 - b. Movie
 - c. Bird

Weapons

1. Knife
 - a. Fork
 - b. Spoon
2. Arrow
 - a. Indian
 - b. Target
3. Bomb
 - a. War
 - b. Death
4. Brick
 - a. Wood
5. Shoe
 - a. Sock
6. Axe
 - a. Tree
 - b. Saw
7. Sword
 - a. Sheath
 - b. Shield
8. Screwdriver
 - a. Screw
 - b. Vodka
 - c. Wrench

Experiment 3

Furniture

1. Sofa
 - a. Sleep
 - b. Sex
 - c. Rest
2. Couch
 - a. Potato
3. Table
 - a. Cloth
 - b. Setting
4. Desk
 - a. Work
 - b. Paper
 - c. Pen
5. Bureau
 - a. Investigation
 - b. Government
 - c. Organization
6. Chest
 - a. Muscle
 - b. Breast
 - c. Heart
7. Vanity
 - a. Beauty
 - b. Pride
 - c. Singer
8. Lounge
 - a. Beer
 - b. Drink
 - c. Liquor
9. Cabinet
 - a. Plate
 - b. Book
10. Bench
 - a. Baseball
 - b. Basketball
 - c. Football
11. Stool
 - a. Feces
 - b. Foot
 - c. Sample
12. Drawer
 - a. Clothes
13. Piano
 - a. Music
 - b. Organ
 - c. Guitar
14. Cushion
 - a. Pin
15. Cupboard
 - a. Cup
 - b. Dishes
 - c. Food
16. Stereo
 - a. Radio
 - b. CD
 - c. Cassette
17. Mirror
 - a. Reflection
18. Television
 - a. Commercial
 - b. Movie
 - c. Video
19. Bar
 - a. Drink
 - b. Alcohol
 - c. Grill
20. Shelf
 - a. Awards
 - b. Life
 - c. Room
21. Rug
 - a. Carpet
 - b. Floor
 - c. Mat
22. Pillow
 - a. Head
 - b. Blanket
 - c. Sheet
23. Radio
 - a. Tape
24. Counter
 - a. Sink
 - b. Clerk
25. Drape
 - a. Window
 - b. House
 - c. Wall
26. Refrigerator
 - a. Microwave
27. Closet
 - a. Hanger
28. Vase
 - a. Flower
 - b. Rose
 - c. Plant
29. Ashtray
 - a. Cigarette
 - b. Ash
 - c. Car
30. Fan
 - a. Air
31. Telephone
 - a. Answering machine

Experiment 3

Vehicles

1. Truck
 - a. Road
 - b. Truck stop
2. Car
 - a. Money
 - b. Race
3. Bus
 - a. Bus stop
 - b. City
 - c. Station
4. Jeep
 - a. Army
 - b. Mud
5. Ambulance
 - a. Hospital
6. Motorcycle
 - a. Helmet
 - b. Mouse
7. Van
 - a. Hippie
 - b. Family
8. Train
 - a. Track
9. Bicycle
 - a. Shop
10. Carriage
 - a. House
11. Airplane
 - a. Airport
 - b. Bird
 - c. Movie
12. Bike
 - a. Rack
 - b. Lock
 - c. Trail
13. Boat
 - a. Water
 - b. Ocean
 - c. Fish
14. Ship
 - a. Sea
 - b. Yard
15. Tractor
 - a. Farm
 - b. Grass
 - c. Dirt
16. Wagon
 - a. Cowboy
17. Subway
 - a. Token
18. Trailer
 - a. Park
19. Cart
 - a. Groceries
 - b. Basket
 - c. Supermarket
20. Yacht
 - a. Club
21. Tank
 - a. Gas
22. Tricycle
 - a. Children
 - b. Baby
23. Canoe
 - a. Indian
 - b. River
24. Raft
 - a. Lake
 - b. Pool
 - c. Beach
25. Submarine
 - a. Sandwich
 - b. Navy
26. Horse
 - a. Cow
 - b. Shit
 - c. Dog
27. Rocket
 - a. Space
 - b. Moon
 - c. Sky
28. Skates
 - a. Ice
 - b. Rink
 - c. Rat
29. Camel
 - a. Desert
 - b. Cigarette
 - c. Egypt
30. Feet
 - a. Shoes
 - b. Hands
 - c. Inch
31. Skis
 - a. Snow
 - b. Mountain
 - c. Resort
32. Elevator
 - a. Stair
 - b. Escalator
 - c. Building

APPENDIX F
RESULTS FROM EXPERIMENTS

Experiment 1

Table F1.
Category Inclusion Rates by Politics and Mode for Female Participants

Group	High	Medium	Atypical	Non-category
Liberal				
Approach	.83(.04)	.84(.04)	.68(.04)	.18(.03)
Neutral	.83(.03)	.84(.04)	.74(.04)	.19(.03)
Avoidance	.88(.03)	.90(.03)	.77(.03)	.21(.02)
Conservative				
Approach	.75(.04)	.74(.04)	.66(.04)	.18(.03)
Neutral	.86(.04)	.87(.04)	.74(.04)	.26(.03)
Avoidance	.87(.05)	.87(.05)	.73(.06)	.12(.04)

Table F2.
Category Inclusion Rates by Politics and Mode for Male Participants

Group	High	Medium	Atypical	Non-category
Liberal				
Approach	.86(.06)	.92(.06)	.82(.06)	.21(.05)
Neutral	.92(.07)	.92(.07)	.65(.77)	.10(.06)
Avoidance	.91(.06)	.94(.06)	.81(.06)	.23(.05)
Conservative				
Approach	.91(.06)	.87(.06)	.72(.06)	.20(.05)
Neutral	.97(.06)	.99(.06)	.71(.07)	.16(.05)
Avoidance	.89(.06)	.91(.06)	.74(.07)	.20(.05)

Table F3.

Median (and IQR) Best Fitting Individual EBRW Parameters for Participants Who Learned the Category Structure

Group	c	w_{range}	A	B	α	μ	k	WSSD
Liberal								
Approach	1.28 (1.64)	0.73 (0.30)	2.75 (41.99)	2.31 (7.29)	132.32 (344.06)	128.34 (6667.11)	0.33 (5.54)	73.38 (95.94)
Neutral	0.76 (1.33)	0.72 (0.37)	5.64 (15.71)	1.51 (10.92)	38.70 (148.14)	50.76 (529.97)	5.62 (15.25)	72.09 (96.82)
Avoidance	1.09 (1.58)	0.76 (0.25)	9.01 (28.89)	4.75 (16.32)	31.06 (73.44)	7.31 (442.74)	2.51 (6.87)	87.38 (121.01)
Conservative								
Approach	1.56 (1.33)	0.67 (0.43)	3.65 (18.85)	3.01 (21.97)	43.03 (89.37)	81.77 (503.18)	4.00 (8.23)	99.57 (73.42)
Neutral	1.28 (1.83)	0.59 (0.17)	6.31 (12.13)	2.89 (322.64)	62.56 (294.25)	408.74 (501.56)	0.54 (4.69)	68.68 (199.23)
Avoidance	1.56 (1.05)	0.66 (0.40)	40.38 (332.25)	2.22 (35.14)	49.66 (128.91)	331.81 (538.09)	0.12 (503.18)	64.80 (182.68)

Note. Parameters c , w_{range} , A , B , and α were each subjected to a separate 3 (Mode) \times 2 (Politics) between subjects ANOVA. Due to the skewed distribution of these parameters (aside from w_{range} which was constrained), outliers in each parameter (as identified by the 1.5/IQR method) were removed and the remaining parameters were logistically transformed. For c , there was no effect of Mode ($F(2,78) = 0.84, p = .46$), Politics ($F(1,78) = 0.99, p = .32$), nor their interaction ($F(2,78) = 0.54, p = .59$). For w_{range} , there was no effect of Mode ($F(2,85) = 0.45, p = .64$), Politics ($F(1,85) = 2.03, p = .16$) nor their interaction ($F(2,72) = 0.91, p = .41$). For A , there was no effect of Mode ($F(2,72) = 0.11, p = .90$), Politics ($F(1,72) = 0.18, p = .68$) nor their interaction ($F(2,70) = 0.42, p = .66$). For B , there was no effect of Mode ($F(2,75) = 0.43, p = .65$), Politics ($F(1,75) = 0.01, p = .91$) nor their interaction ($F(2,75) = 1.05, p = .36$). For α , there was no effect of Mode ($F(1,74) = 0.11, p = .89$), Politics ($F(1,74) = 0.49, p = .49$) nor their interaction ($F(2,74) = 0.79, p = .46$). Finally, for k , there was no effect of Mode ($F(1,73) = 0.27, p = .76$), Politics ($F(1,73) = 1.16, p = .28$) nor their interaction ($F(2,73) = 1.81, p = .17$).

Table F4.

Best Fitting Parameters and Mean Cross-Validation WSSD Values for Averaged Data

	c	w_{range}	A	B	α	k	μ	WSSD
C	0.36	0.45	27.37	3.78	578.98	0.007	670.69	410.10
Liberals	0.65	0.45	5.35	2.03	107 × 3.30	10-6 × 1.66	557.64	220.29
Conservatives	0.92							
Approach	0.48	0.43	7.24	2.32	108 × 3.90	10-8 × 9.10	560.51	122.29
Avoidance	0.84							
Conservative	1.20	0.45	6.89	2.40	108 × 9.00	10-8 × 4.15	532.52	98.75
Avoidance								
Others	0.58							
W								
Liberal	0.30	0.23	55.68	5.32	1880.81	0.0008	659.86	397.80
Conservative		0.47						
Approach	0.22	0.21	41.41	6.20	1149.76	0.0017	628.78	380.82
Avoidance		0.51						
Conservative	0.24	0.51	43.31	6.27	216.89	0.008	584.66	378.29
Avoidance		0.24						
Others		0.24						
A								
Liberal	0.48	0.38	13.80	2.73	105 × 6.22	10-5 × 2.16	629.16	376.25
Conservative			11.93					
Approach	0.51	0.40	9.31	2.56	107 × 5.01	10-7 × 3.72	615.91	346.07
Avoidance			11.96					
Conservative	0.37	0.39	27.02	3.61	107 × 3.91	10-7 × 1.36	602.11	349.64
Avoidance			19.30					
Others								
B								
Liberal	0.62	0.40	6.78	2.39	105 × 1.75	10-4 × 2.11	589.66	314.95
Conservative				1.83				
Approach	0.44	0.42	17.21	3.05	104 × 2.31	10-4 × 3.77	654.03	275.91
Avoidance				3.11				
Conservative	0.83	0.44	3.92	1.30	106 × 2.45	10-5 × 4.32	395.18	349.03
Avoidance				1.97				
Others								

Table F4 continued

<i>A & B</i>	<i>c</i>	<i>W_{range}</i>	<i>A</i>	<i>B</i>	α	<i>k</i>	μ	<i>WSSD</i>
Liberal	0.63	0.40	6.49	2.42	105 × 1.73	10-4 × 2.25	582.41	334.93
Conservative			6.75	1.78				
Approach	0.76	0.44	5.00	2.07	106 × 3.73	10-5 × 1.86	514.48	390.80
Avoidance			4.46	1.74				
Conservative			4.89	1.53	105 × 1.74	10-4 × 3.65	454.95	431.57
Avoidance	0.71	0.42						
Others			5.37	2.20				

Experiment 2

Table F5.
Best Fitting Parameters and AIC Values for Clothing

	<i>c</i>	<i>W_{range}</i>	<i>Y</i>	<i>AIC</i>
Constrained	31.63	0.80	4.26	87.79
<i>c</i>				
Liberals	31.12			
Conservatives	32.06	0.80	4.26	89.63
Approach	31.88			
Avoidance	31.43	0.80	4.26	89.76
Conservative				
Avoidance	34.90	0.80	4.22	87.64
Others	30.55			
<i>W_{var}</i>				
Liberal		0.81		
Conservative	31.78	0.80	4.26	93.12
Approach		0.78		
Avoidance	32.13	0.81	4.26	91.32
Conservative				
Avoidance	31.79	0.69	4.14	92.05
Others		0.81		
<i>Y</i>				
Liberal			4.19	
Conservative	31.63	0.80	4.33	89.72
Approach			5.06	
Avoidance	31.08	0.80	4.03	86.85
Conservative			4.31	
Avoidance	31.67	0.80		89.76
Others			4.23	

Table F6.
Best Fitting Parameters and AIC Values for Furniture

	<i>c</i>	<i>W</i> _{range}	<i>Y</i>	<i>AIC</i>
Constrained	19.25	0.46	3.90	215.31
<i>c</i>				
Liberals	20.23	0.51	4.10	216.19
Conservatives	18.59			
Approach	21.10	0.53	3.63	225.93
Avoidance	21.59			
Conservative	18.94	0.44	3.94	222.99
Avoidance				
Others	19.94			
<i>W</i> _{var}				
Liberal	23.38	0.45	3.43	243.41
Conservative		0.52		
Approach	18.07	0.63	4.36	241.91
Avoidance		0.31		
Conservative	19.82	0.44	3.74	241.22
Avoidance				
Others		0.47		
<i>Y</i>				
Liberal	20.14	0.51	4.33	215.01
Conservative			3.58	
Approach	17.72	0.39	3.99	218.24
Avoidance			4.41	
Conservative	17.82	0.38	4.49	227.39
Avoidance				
Others			4.07	

Table F7.
Best Fitting Parameters and AIC Values for Vehicles

	<i>c</i>	<i>W</i> _{range}	<i>Y</i>	<i>AIC</i>
Constrained	6.33	0.60	11.98	311.92
<i>c</i>				
Liberals	6.81			
Conservatives	6.39	0.62	12.07	314.96
Approach	6.94			
Avoidance	6.32	0.61	11.77	314.35
Conservative				
Avoidance	5.92	0.60	12.18	317.20
Others	6.59			
<i>W</i> _{var}				
Liberal		0.61		
Conservative	6.63	0.63	11.73	333.42
Approach		0.63		
Avoidance	7.96	0.57	9.27	334.46
Conservative				
Avoidance	6.48	0.68	11.08	335.12
Others		0.51		
<i>Y</i>				
Liberal			11.82	
Conservative	6.91	0.59	10.25	315.06
Approach			11.78	
Avoidance	6.87	0.60	10.42	316.57
Conservative				
Avoidance	8.05	0.52	9.97	304.95
Others			8.05	

Experiment 3

Table F8.
Typicality, Familiarity and “In-The-Category” Judgments for Furniture

Stimulus	Norming Study		Categorization Study				Liberal Avoid	Conservative Approach	Conservative Neutral	Conservative Avoid
	Typicality	Familiarity	Liberal Approach	Liberal Neutral	Liberal Avoid	Conservative Approach				
chair	8.98	8.98	1.00	1.00	1.00	1.00	0.94	1.00	1.00	
sofa	8.94	8.98	1.00	1.00	0.95	0.95	1.00	1.00	1.00	
couch	9.00	8.98	1.00	1.00	0.95	0.95	0.94	0.95	0.89	
table	9.00	8.96	0.95	1.00	1.00	1.00	1.00	0.95	1.00	
easy chair	6.89	7.17	1.00	0.95	1.00	1.00	0.94	1.00	1.00	
dresser	8.32	8.87	0.95	1.00	1.00	1.00	1.00	0.95	1.00	
rocking chair	8.21	8.79	0.95	1.00	1.00	1.00	0.88	1.00	0.94	
coffee table	8.70	8.96	0.95	0.95	1.00	1.00	1.00	1.00	1.00	
rocker	7.36	8.38	1.00	1.00	0.95	0.95	0.76	0.89	0.94	
love seat	8.04	8.26	1.00	0.95	1.00	1.00	1.00	1.00	0.94	
chest of drawers	7.51	8.15	0.95	1.00	1.00	1.00	0.94	0.95	1.00	
desk	8.60	8.94	0.95	0.95	1.00	1.00	1.00	1.00	1.00	
bed	8.47	8.96	1.00	1.00	1.00	1.00	0.94	0.95	1.00	
bureau	8.00	8.47	0.95	0.95	1.00	1.00	0.65	0.95	0.94	
davenport	2.87	2.55	0.52	0.43	0.45	0.45	0.18	0.26	0.44	
divan	3.47	3.15	0.57	0.29	0.30	0.30	0.47	0.16	0.33	
chest	7.68	8.40	0.95	0.90	0.95	0.95	0.88	0.79	1.00	
cedar chest	6.74	6.81	0.95	0.90	0.95	0.95	0.88	0.84	0.94	
vanity	6.28	7.11	0.90	0.71	0.90	0.90	0.76	0.74	0.89	

Table F8 continued

Stimulus	Norming Study			Categorization Study					
	Typicality	Familiarity		Liberal Approach	Liberal Neutral	Liberal Avoid	Conservative Approach	Conservative Neutral	Conservative Avoid
bookcase	7.70	8.64		0.95	0.95	0.90	0.82	0.95	0.94
lounge	5.11	8.02		0.90	0.76	0.75	0.76	0.58	0.83
chaise lounge	4.91	5.55		0.71	0.76	0.90	0.65	0.47	0.83
ottoman	6.64	7.13		0.90	0.95	0.90	0.76	0.84	0.89
footstool	6.72	8.62		0.95	0.95	1.00	0.82	0.95	1.00
cabinet	6.83	8.89		0.76	0.95	0.85	0.71	0.74	0.89
china closet	6.47	7.79		0.76	0.71	0.95	0.71	0.79	0.89
bench	7.32	8.94		0.95	1.00	0.95	0.94	0.89	1.00
buffet	2.91	7.04		0.33	0.33	0.30	0.24	0.11	0.28
lamp	6.11	8.72		0.43	0.24	0.40	0.29	0.53	0.56
stool	7.98	8.89		0.95	1.00	0.95	0.82	0.89	1.00
hassock	2.60	2.30		0.29	0.52	0.35	0.29	0.21	0.50
drawers	7.17	8.85		0.86	0.90	0.80	0.82	0.74	0.94
piano	4.89	8.81		0.38	0.33	0.40	0.24	0.47	0.50
cushion	6.00	8.72		0.57	0.43	0.45	0.47	0.53	0.33
magazine rack	4.83	8.13		0.43	0.29	0.50	0.53	0.63	0.61
hi-fi	2.79	6.57		0.10	0.05	0.10	0.00	0.11	0.17
cupboard	5.96	8.57		0.81	0.71	0.65	0.53	0.63	0.78
stereo	3.57	8.79		0.05	0.05	0.05	0.00	0.11	0.11
mirror	5.83	8.89		0.29	0.33	0.40	0.47	0.53	0.28
television	4.85	8.89		0.00	0.24	0.25	0.18	0.42	0.39
bar	4.72	8.64		0.52	0.43	0.65	0.59	0.63	0.78

Table F8 continued

Stimulus	Norming Study		Categorization Study					
	Typicality	Familiarity	Liberal Approach	Liberal Neutral	Liberal Avoid	Conservative Approach	Conservative Neutral	Conservative Avoid
shelf	7.13	8.81	0.67	0.71	0.70	0.59	0.68	1.00
rug	5.51	8.94	0.24	0.24	0.40	0.35	0.37	0.39
pillow	4.06	8.91	0.19	0.00	0.10	0.12	0.16	0.17
wastebasket	4.32	8.81	0.24	0.14	0.45	0.24	0.37	0.39
radio	3.21	8.87	0.05	0.00	0.00	0.00	0.00	0.06
sewing machine	2.60	8.49	0.10	0.00	0.10	0.00	0.05	0.11
stove	4.55	8.98	0.33	0.29	0.15	0.18	0.42	0.44
counter	5.87	8.94	0.57	0.43	0.50	0.35	0.47	0.61
clock	4.70	8.96	0.14	0.10	0.15	0.18	0.26	0.22
drapes	4.32	7.98	0.24	0.05	0.15	0.12	0.16	0.17
refrigerator	4.11	8.96	0.29	0.33	0.25	0.24	0.32	0.39
picture	3.85	8.96	0.05	0.00	0.10	0.00	0.16	0.00
closet	4.62	8.98	0.43	0.29	0.50	0.29	0.42	0.67
vase	4.06	8.70	0.10	0.00	0.25	0.00	0.21	0.22
ashtray	2.28	8.26	0.00	0.05	0.00	0.06	0.16	0.11
fan	3.83	8.79	0.00	0.00	0.05	0.06	0.16	0.00
telephone	3.23	8.96	0.00	0.00	0.05	0.12	0.05	0.11

Table F9.

Typicality, Familiarity and “In-The-Category” Judgments for Vehicles

Stimulus	Norming Study		Categorization Study				Liberal Avoid	Conservative Approach	Conservative Neutral	Conservative Avoid
	Typicality	Familiarity	Liberal Approach	Liberal Neutral	Liberal Avoid	Conservative Approach				
automobile	8.98	8.96	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
truck	8.85	8.94	1.00	0.95	1.00	1.00	0.94	0.89	1.00	
car	9.00	8.98	1.00	1.00	0.95	1.00	1.00	1.00	1.00	
bus	8.79	8.87	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
taxi	8.55	8.79	1.00	0.95	1.00	1.00	1.00	0.95	1.00	
jeep	8.62	8.96	1.00	1.00	1.00	1.00	1.00	0.95	1.00	
ambulance	7.91	8.83	1.00	1.00	0.95	0.94	0.94	0.95	1.00	
motorcycle	8.04	8.91	1.00	0.95	1.00	1.00	1.00	1.00	0.94	
streetcar	7.30	7.72	1.00	1.00	0.95	0.94	0.94	0.95	0.94	
van	8.57	8.96	1.00	1.00	1.00	1.00	1.00	0.95	1.00	
train	7.45	8.85	0.95	1.00	0.85	0.94	0.94	0.79	0.83	
trolley	5.62	7.96	0.95	0.81	0.80	0.76	0.76	0.79	1.00	
bicycle	5.62	8.89	0.77	0.76	0.60	0.71	0.71	0.68	0.78	
carriage	3.89	8.28	0.73	0.33	0.65	0.53	0.53	0.58	0.50	
airplane	6.70	8.96	0.86	0.95	0.65	0.76	0.76	0.74	0.78	
bike	6.00	8.85	0.77	0.62	0.60	0.53	0.53	0.63	0.67	
boat	5.87	8.70	0.73	0.71	0.55	0.71	0.71	0.68	0.67	
jet	5.40	8.55	0.73	0.81	0.65	0.82	0.82	0.63	0.72	
ship	5.45	8.60	0.64	0.71	0.65	0.71	0.71	0.58	0.78	
scooter	4.55	8.57	0.68	0.38	0.45	0.53	0.53	0.42	0.67	
tractor	5.47	8.40	0.86	0.90	0.80	0.76	0.76	0.79	0.94	
wagon	4.04	8.26	0.73	0.33	0.40	0.35	0.35	0.47	0.50	

Table F9 continued

Stimulus	Norming Study			Categorization Study					
	Typicality	Familiarity		Liberal Approach	Liberal Neutral	Liberal Avoid	Conservative Approach	Conservative Neutral	Conservative Avoid
subway	6.51	8.77		0.91	0.86	0.80	0.82	0.89	0.72
trailer	5.72	8.51		0.73	0.62	0.55	0.76	0.63	0.61
cart	3.40	8.09		0.45	0.33	0.30	0.35	0.47	0.50
wheelchair	3.30	8.66		0.50	0.29	0.25	0.29	0.26	0.56
yacht	4.77	7.85		0.68	0.67	0.60	0.65	0.58	0.72
tank	4.53	8.06		0.73	0.67	0.65	0.88	0.68	0.83
go-cart	4.70	8.02		0.77	0.67	0.55	0.76	0.79	0.83
rowboat	4.15	8.23		0.50	0.38	0.40	0.29	0.47	0.39
dogsled	2.94	7.60		0.55	0.33	0.30	0.18	0.32	0.33
tricycle	3.53	8.51		0.68	0.62	0.45	0.53	0.42	0.56
canoe	3.72	8.53		0.36	0.38	0.25	0.24	0.37	0.39
raft	3.62	7.94		0.32	0.29	0.25	0.18	0.32	0.22
submarine	3.81	7.98		0.73	0.76	0.60	0.59	0.58	0.72
sled	2.45	8.68		0.41	0.29	0.15	0.24	0.26	0.17
horse	3.38	8.62		0.45	0.19	0.25	0.12	0.16	0.22
rocket	2.91	8.06		0.73	0.48	0.45	0.41	0.53	0.44
blimp	3.26	7.21		0.55	0.38	0.45	0.41	0.47	0.67
skates	2.72	8.68		0.36	0.24	0.10	0.06	0.16	0.33
camel	2.47	7.89		0.32	0.14	0.20	0.06	0.16	0.11
feet	3.28	8.81		0.18	0.05	0.10	0.06	0.11	0.11
skis	2.34	8.47		0.32	0.29	0.05	0.18	0.16	0.11

Table F9 continued

Stimulus	Norming Study		Categorization Study					
	Typicality	Familiarity	Liberal Approach	Liberal Neutral	Liberal Avoid	Conservative Approach	Conservative Neutral	Conservative Avoid
skateboard	3.21	8.30	0.41	0.19	0.15	0.24	0.32	0.28
wheelbarrow	2.19	8.26	0.27	0.14	0.20	0.18	0.00	0.33
surfboard	2.47	8.13	0.18	0.10	0.00	0.12	0.05	0.06
elevator	2.72	8.83	0.32	0.19	0.20	0.18	0.16	0.22

Table F10.
Best Fitting Parameters and WSSD Values for Furniture

	c	w_{range}	A	B	α	k	μ	WSSD
Constrained	52.58	0.56	1.03	1.76	44.58	13.09	2.04	50.93
c								
Liberals	44.94	0.48	1.40	2.71	0.78	144.4	437.40	49.98
Conservatives	38.31							
Approach	58.91	0.77	1.02	1.77	25.95	22.13	12.77	54.73
Avoidance	57.10							
Conservative	53.33	0.79	1.00	1.63	133.19	4.98	3.18	48.95
Avoidance	61.59							
Liberals	54.71	0.75	1.00	1.74	39.68	15.10	5.53	53.18
Approach	48.95							
w								
Liberal	48.25	0.78	1.01	1.76	256.63	2.33	7.12	48.50
Conservative		0.70						
Approach	51.63	0.84	1.08	1.85	31.65	17.09	2.25	52.15
Avoidance		0.50						
Conservative	44.44	0.52	1.03	1.83	108.80	5.11	16.11	49.07
Avoidance		0.76						
Liberals	50.88	0.58	1.03	1.77	65.79	8.91	3.62	51.78
Approach	50.88							
Liberals	50.88	0.58	1.03	1.77	65.79	8.91	3.62	51.78
Approach		0.65						
A								
Liberal	44.21	0.74	1.08	1.86	8.09	59.46	91.84	53.95
Conservative			1.07					
Approach	44.28	0.62	1.03	1.84	2.76	154.62	61.95	50.49
Avoidance			1.04					
Conservative	51.17	0.67	0.94	1.73	39.56	15.19	7.41	48.35
Avoidance			1.04					
Liberals	53.87	0.62	0.93	1.71	40.78	14.84	15.12	51.10
Approach			1.02					
B								
Liberal	73.63	0.90	0.94	1.54	98.21	6.96	0.58	52.61
Conservative				1.64				
Approach	53.76	0.56	1.08	1.80	6.12	87.82	-22.83	52.75
Avoidance				1.75				
Conservative	55.66	0.67	0.96	1.54	204.15	3.26	3.88	54.76
Avoidance				1.67				
Liberals	42.65	0.65	1.05	1.79	96.12	5.69	6.02	51.89
Approach				1.89				
Liberals	42.65	0.65	1.05	1.79	96.12	5.69	6.02	51.89
Approach				1.89				

Table F10 continued.

	<i>c</i>	<i>w</i> _{range}	<i>A</i>	<i>B</i>	α	<i>k</i>	μ	<i>WSSD</i>
<i>A & B</i>								
Liberal	55.54	0.72	0.99	1.79	101.27	6.38	2.23	51.81
Conservative			1.03	1.77				
Approach	64.70	0.74	0.96	1.68	137.66	4.30	0.39	48.56
Avoidance			0.97	1.62				
Conservative	39.05	0.65	1.05	1.89	3.09	131.82	78.74	49.52
Avoidance			1.13	1.97				
Others								
Liberals	49.05	0.59	0.91	1.90	29.99	19.48	-0.07	110.50
Approach			1.07	1.72				
Others								

Table F11.
Best Fitting Parameters and WSSD Values for Vehicles

	c	w_{range}	A	B	α	k	μ	WSSD
Constrained	3.75	0.43	28.59	38.10	1.82	0.56	-24.50	658.45
c								
Liberals	2.18	0.52	56.12	45.43	4.63	0.09	2.36	579.98
Conservatives	2.11							
Approach	3.63	0.49	23.96	32.83	1.98	0.70	2.24	490.49
Avoidance	3.81							
Conservative	3.68	0.50	28.13	33.16	6.72	0.17	1.63	548.00
Avoidance								
Others	3.34							
Liberals	2.95	0.54	34.97	36.17	2.47	0.35	-2.15	397.11
Approach								
Others	2.75							
w								
Liberal	3.49	0.54	23.80	28.94	5.72	0.28	-14.51	547.67
Conservative		0.52						
Approach	3.64	0.53	23.92	28.18	15.58	0.10	-3.52	459.08
Avoidance		0.50						
Conservative	2.80	0.41	38.30	39.90	8.93	0.81	2.88	574.15
Avoidance								
Others		0.54						
Liberals	3.15	0.55	26.48	31.16	3.31	0.40	-7.30	544.40
Approach								
Others		0.51						
A								
Liberal	3.20	0.49	31.94	33.79	3.65	0.24	101.04	505.58
Conservative			31.30					
Approach	1.52	0.54	170.86	57.31	7.49	0.01	416.96	734.00
Avoidance			153.72					
Conservative	3.77	0.40	28.79	34.00	21.39	0.05	-5.79	512.16
Avoidance								
Others			32.66					
Liberals	3.55	0.52	22.65	29.56	0.66	2.20	-5.45	549.09
Approach								
Others			25.19					
B								
Liberal	5.01	0.41	20.36	25.31	8.12	0.27	-9.62	472.20
Conservative				24.42				
Approach	4.46	0.49	19.61	28.95	34.74	0.06	-20.81	468.76
Avoidance				25.27				
Conservative	4.12	0.44	23.20	27.44	14.48	0.11	-5.35	404.35
Avoidance								
Others				31.57				
Liberals	3.23	0.51	27.94	34.78	2.25	0.52	-9.83	440.99
Approach								
Others				32.71				

Table F11 continued.

	<i>c</i>	<i>w_{range}</i>	<i>A</i>	<i>B</i>	<i>α</i>	<i>k</i>	<i>μ</i>	<i>WSSD</i>
<i>A & B</i>								
Liberal	4.10	0.44	21.44	32.05	2.35	0.68	-12.27	363.17
Conservative			23.09	29.51				
Approach	3.69	0.49	22.45	39.22	4.67	0.29	-1.50	396.18
Avoidance			23.29	32.77				
Conservative	3.99	0.49	20.12	28.06	3.30	0.54	-2.62	381.26
Avoidance			22.85	28.04				
Others								
<i>Liberals</i>								
Approach	3.77	0.48	23.43	34.96	4.68	0.29	-2.05	463.80
Others			26.04	31.20				

REFERENCES

- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409-429.
- Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, *37*(3), 372-400.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*(3), 211-227.
- Blair, M., & Homa, D. (2003). As easy to memorize as they are to classify: The 5-4 categories and the category advantage. *Memory & Cognition*, *31*(8), 1293-1301.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433-436.
- Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, *120*(3), 278-287.
- Browne, M. W. (2000). Cross-validation methods. *Journal of Mathematical Psychology*, *44*(1), 108-132.
- Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods, Instruments & Computers*, *41*(4), 977-990.
- Budanitsky, A., & Hirst, G. (2006). Evaluating WordNet-based measures of lexical semantic relatedness. *Computational Linguistics*, *32*(1), 13-47.
- Burnett, R. C., Medin, D. L., Ross, N. O., & Blok, S. V. (2005). Ideal is typical. *Canadian Journal of Experimental Psychology*, *59*(1), 3-10.
- Burnham, K. P., & Anderson, D. (2002). *Model selection and multi-model inference: A practical information-theoretic approach* (2nd ed.). New York: Springer-Verlag.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning & Verbal Behavior*, *8*(2), 240-247.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, *41*(6), 391-407.
- Dry, M. J., & Storms, G. (2009). Similar but not the same: A comparison of the utility of directly rated and feature-based similarity measures for generating spatial models of conceptual data. *Behavior Research Methods*, *41*(3), 889-900.

- Ell, S. W., Cosley, B., & McCoy, S. K. (2011). When bad stress goes good: Increased threat reactivity predicts improved category learning performance. *Psychonomic Bulletin & Review*, *18*(1), 96-102.
- Emmanuel, K., Kevin, D., & Marc, B. (2010). Practice effects in large-scale visual word recognition studies: A lexical decision study on 14,000 dutch mono- and disyllabic words and nonwords. *Frontiers in Psychology*, *1*(November), 1-15.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*(2), 107-140.
- Förster, J. (2009). Relations between perceptual and conceptual scope: How global versus local processing fits a focus on similarity versus dissimilarity. *Journal of Experimental Psychology: General*, *138*(1), 88-111.
- Förster, J., Friedman, R. S., Özelsel, A., & Denzler, M. (2006). Enactment of approach and avoidance behavior influences the scope of perceptual and conceptual attention. *Journal of Experimental Social Psychology*, *42*(2), 133-146.
- Friedman, R. S., & Förster, J. (2001). The effects of promotion and prevention cues on creativity. *Journal of Personality and Social Psychology*, *81*(6), 1001-1013.
- Friedman, R. S., & Förster, J. (2005). The influence of approach and avoidance cues on attentional flexibility. *Motivation and Emotion*, *29*(2), 69-81.
- Gable, P. A., & Harmon-Jones, E. (2008). Approach-motivated positive affect reduces breadth of attention. *Psychological Science*, *19*(5), 476-482.
- Gable, P., & Harmon-Jones, E. (2010). The blues broaden, but the nasty narrows: Attentional consequences of negative affects low and high in motivational intensity. *Psychological Science*, *21*(2), 211-215.
- Hampton, J. A. (1993). Prototype models of concept representation. In I. van Mechelen, J. A. Hampton, R. S. Michalski, & P. Theuns (Ed.), *Categories and concepts: Theoretical views and inductive data analysis*. (pp. 67-95)
- Heit, E., Briggs, J., & Bott, L. (2004). Modeling the effects of prior knowledge on learning incongruent features of category members. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(5), 1065-1081.
- Hoffman, A. B., Harris, H. D., & Murphy, G. L. (2008). Prior knowledge enhances the category dimensionality effect. *Memory & Cognition*, *36*(2), 256-270.
- Isen, A. M., & Daubman, K. A. (1984). The influence of affect on categorization. *Journal of Personality and Social Psychology*, *47*(6), 1206-1217.

- Janoff-Bulman, R., Sheikh, S., & Baldacci, K. G. (2008). Mapping moral motives: Approach, avoidance, and political orientation. *Journal of Experimental Social Psychology, 44*(4), 1091-1099.
- Kaplan, A. S., & Murphy, G. L. (1999). The acquisition of category structure in unsupervised learning. *Memory & Cognition, 27*(4), 699-712.
- Kaplan, A. S., & Murphy, G. L. (2000). Category learning with minimal prior knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(4), 829-846.
- Kelemen, D., & Bloom, P. (1994). Domain-specific knowledge in simple categorization tasks. *Psychonomic Bulletin & Review, 1*(3), 390-395.
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in psychtoolbox-3? *Perception 36*, ECVF Abstract Supplement.
- Koch, S., Holland, R. W., Hengstler, M., & van Knippenberg, A. (2009). Body locomotion as regulatory process: Stepping backward enhances cognitive control. *Psychological Science, 20*(5), 549-550.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*(1), 22-44.
- Kruschke, J. K. (2008). Models of categorization. In R. Sun, & R. Sun (Eds.), *The Cambridge handbook of computational psychology*. (pp. 267-301). New York, NY US: Cambridge University Press.
- Lamberts, K. (1995). Categorization under time pressure. *Journal of Experimental Psychology: General, 124*(12), 161-180.
- Lamberts, K. (1998). The time course of categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*(3), 695-711.
- Landauer, T. K., Foltz, P. W., & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes, 25*(2-3), 259-284.
- Lin, E. L., & Murphy, G. L. (1997). Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception and Performance, 23*(4), 1153-1169.
- Lombrozo, T. (2009). Explanation and categorization: How 'why?' informs 'what?'. *Cognition, 110*(2), 248-253.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111*(2), 309-332.

- Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory & Cognition, 28*(1), 41-50.
- 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), 397-407.
- Maddox, W. T. (1999). On the dangers of averaging across observers when comparing decision bound models and generalized context models of categorization. *Perception & Psychophysics, 61*(2), 354-375.
- Maddox, W. T., & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics, 53*(1), 49-70.
- Markman, A. B., Baldwin, G. C., & Maddox, W. T. (2005). The interaction of payoff structure and regulatory focus in classification. *Psychological Science, 16*(11), 852-855.
- McCloskey, M. E., & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory & Cognition, 6*(4), 462-472.
- Medin, D. L., Cox, D. G., & Ross, N. O. (2006). *Culture and resource conflict: Why meanings matter*. New York, NY US: Russell Sage Foundation.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*(3), 207-238.
- Minda, J. P., & Smith, J. D. (2002). Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(2), 275-292.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Murphy, G. L., & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(4), 904-919.
- Murphy, G. L., & Medin, D. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*(3), 289-316.
- Murphy, G. L., & Wisniewski, E. J. (1989). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(4), 572-586.
- Nash, K., McGregor, I., & Inzlicht, M. (2010). Line bisection as a neural marker of approach motivation. *Psychophysiology, 47*(5), 979-983.

- Navarro, D. J. (2007). On the interaction between exemplar-based concepts and a response scaling process. *Journal of Mathematical Psychology*, 51(2), 85-98.
- Nelson, D. L., McEvoy, C. L. & Schreiber, T. A. (1998). *The University of South Florida word association, rhyme, and word fragment norms*. Retrieved February, 2011, from <http://w3.usf.edu/FreeAssociation/>
- Nix v. Hedden*, 149 U.S. 304 (1893).
- Nosofsky, R. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1), 104-114.
- Nosofsky, R. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39-57.
- Nosofsky, R. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(1), 87-108.
- Nosofsky, R. (1988). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(1), 54-65.
- Nosofsky, R., & Zaki, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*, 9(4), 247-255.
- Nosofsky, R., & Alfonso-Reese, L. (1999). Effects of similarity and practice on speeded classification response times and accuracies: Further tests of an exemplar-retrieval model. *Memory & Cognition*, 27(1), 78-93.
- Nosofsky, R., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2), 266-300.
- Nosofsky, R., & Stanton, R. D. (2005). Speeded classification in a probabilistic category structure: Contrasting exemplar-retrieval, decision-boundary, and prototype models. *Journal of Experimental Psychology: Human Perception and Performance*, 31(3), 608-629.
- Patalano, A. L., Chin-Parker, S., & Ross, B. H. (2006). The importance of being coherent: Category coherence, cross-classification, and reasoning. *Journal of Memory and Language*, 54(3), 407-424.
- Patalano, A. L., Wengrovitz, S. M., & Sharpes, K. M. (2009). The influence of category coherence on inference about cross-classified entities. *Memory & Cognition*, 37(1), 21-28.

- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437-442.
- Plato. (1975). *The trial & death of Socrates: Plato's Euthyphro, Apology, Crito, and death scene from Phaedo [Dialogues.]*. Indianapolis: Hackett Pub. Co.
- Price, T. F., & Harmon-Jones, E. (2010). The effect of embodied emotive states on cognitive categorization. *Emotion*, 10(6), 934-938.
- Princeton University. (2010). *About WordNet*. Retrieved February, 2011, from <http://wordnet.princeton.edu>
- Ratneshwar, S., Barsalou, L. W., Pechmann, C., & Moore, M. (2001). Goal-derived categories: The role of personal and situational goals in category representations. *Journal of Consumer Psychology*, 10(3), 147-157.
- Rehder, B., & Hoffman, A. B. (2005). Thirty-something categorization results explained: Selective attention, eyetracking, and models of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(5), 811-829.
- Rehder, B., & Murphy, G. L. (2003). A knowledge-resonance (KRES) model of category learning. *Psychonomic Bulletin & Review*, 10(4), 759-784.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning & Verbal Behavior*, 14(6), 665-681.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12(1), 1-20.
- Rock, M., & Janoff-Bulman, R. (2010). Where do we draw our lines? politics, rigidity, and the role of self-regulation. *Social Psychological and Personality Science*, 1(1), 26-33.
- Rodd, J. M. (2004). When do leotards get their spots? Semantic activation of lexical neighbors in visual word recognition. *Psychonomic Bulletin & Review*, 11(3), 434-439.
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104(3), 192-233.
- Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382-439.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4), 573-605.

- Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15(3), 346-378.
- Smith, J. D., & Minda, J. P. (2002). Distinguishing prototype-based and exemplar-based processes in dot-pattern category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(4), 800-811.
- Smits, T., Storms, G., Rosseel, Y., & De Boeck, P. (2002). Fruits and vegetables categorized: An application of the generalized context model. *Psychonomic Bulletin & Review*, 9(4), 836-844.
- Spalding, T. L., & Murphy, G. L. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(2), 525-538.
- Spalding, T. L., & Murphy, G. L. (1999). What is learned in knowledge-related categories? evidence from typicality and feature frequency judgments. *Memory & Cognition*, 27(5), 856-867.
- Storms, G., De Boeck, P., & Ruts, W. (2001). Categorization of novel stimuli in well-known natural concepts: A case study. *Psychonomic Bulletin & Review*, 8(2), 377-384.
- Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, 23(3), 457-482.
- Vanpaemel, W., & Storms, G. (2008). In search of abstraction: The varying-abstraction model of categorization. *Psychonomic Bulletin & Review*, 15(4), 732-749.
- Vanpaemel, W., Verbeemen, T., Dry, M., Verguts, T., & Storms, G. (2010). Geometric and featural representations in semantic concepts. *Memory & Cognition*, 38(7), 962-968.
- Verbeemen, T., Vanpaemel, W., Pattyn, S., Storms, G., & Verguts, T. (2007). Beyond exemplars and prototypes as memory representations of natural concepts: A clustering approach. *Journal of Memory and Language*, 56(4), 537-554.
- Verheyen, S., Ameel, E., & Storms, G. (2007). Determining the dimensionality in spatial representations of semantic concepts. *Behavior Research Methods*, 39(3), 427-438.
- Verheyen, S., De Deyne, S., Dry, M. J., & Storms, G. (2011). Uncovering contrast categories in categorization with a probabilistic threshold model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(6), 1515-1531.

- Verheyen, S., Hampton, J. A., & Storms, G. (2010). A probabilistic threshold model: Analyzing semantic categorization data with the rasch model. *Acta Psychologica, 135*(2), 216-225.
- Verheyen, S., Heussen, D., & Storms, G. (2011). On domain differences in categorization and context variety. *Memory & Cognition, 39*(7), 1290-1300.
- Voorspoels, W., Vanpaemel, W., & Storms, G. (2008). Exemplars and prototypes in natural language concepts: A typicality-based evaluation. *Psychonomic Bulletin & Review, 15*(3), 630-637.
- Wickens, T. D. (1982). *Models for behavior: Stochastic processes in psychology*. San Francisco, CA: W. H. Freeman and Co.
- Wills, A. J., & Pothos, E. M. (2011). On the adequacy of current empirical evaluations of formal models of categorization. *Psychological Bulletin, 138*(1), 102-125.
- Wittgenstein, L. (1953). *Philosophical investigations* (G.E.M. Anscombe Trans.). Oxford: Blackwell.
- Wu, L., & Barsalou, L. W. (2009). Perceptual simulation in conceptual combination: Evidence from property generation. *Acta Psychologica, 132*(2), 173-189.
- Yamauchi, T., & Yu, N. (2008). Category labels versus feature labels: Category labels polarize inferential predictions. *Memory & Cognition, 36*(3), 544-553.
- Zaki, S. R., Nosofsky, R., Stanton, R. D., & Cohen, A. L. (2003). Prototype and exemplar accounts of category learning and attentional allocation: A reassessment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*(6), 1160-1173.
- Zivot, M., Cohen, A. L., & Kapucu, A. (2012). Modeling the effect of mood on dimensional attention during categorization. Manuscript submitted for publication.