

Durham E-Theses

Conceptual centrality and property induction

Constantinos Hadjichristidis

How to cite:

Hadjichristidis, Constantinos (2000) Conceptual centrality and property induction. Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/4339/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

CONCEPTUAL CENTRALITY AND PROPERTY INDUCTION

CONSTANTINOS HADJICHRISTIDIS

Submitted for the qualification of

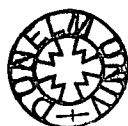
Doctor of Philosophy,

University of Durham,

Psychology Department,

June 2000

The copyright of this thesis rests with the author. No quotation from it should be published in any form, including Electronic and the Internet, without the author's prior written consent. All information derived from this thesis must be acknowledged appropriately.



19 JUN 2001

ABSTRACT

This thesis examines property generalization among concepts. Its primary objective is to investigate the hypothesis that the more central a feature for a concept, the higher its generalizability to other concepts that share a similar structure (features and dependencies). Its secondary objectives are to examine the relative contributions of feature centrality and feature variability in property induction, whether centrality offers a domain-general or a domain-specific constraint, and whether centrality can operate under conditions of vagueness.

Experiments 1 and 2 addressed the centrality hypothesis with centrality measured, whereas Experiments 3 to 14 and 17 with centrality manipulated. Relative feature centrality was manipulated as follows: from a single-dependency chain (Experiments 3 to 7), from the number of properties that depended upon a feature (Experiments 8 to 11 and 17), and from the centrality of the properties that depended upon the critical features (Experiments 12 to 14). The results support the centrality hypothesis.

Experiments 12 to 16 addressed the relative contributions of centrality and variability in property induction. Experiments 12 to 14 pitted a central and variable property against a less central and less variable property in judgments of frequency and inductive strength. The results suggest that property induction depends on centrality rather than frequency information, and that centrality can bias the perception of frequency (although the latter results were not clear-cut). Experiments 15 and 16 pitted centrality against variability in information seeking. The results show that centrality information is sought more often than variability information to make an inference, especially amongst dissimilar concepts.

Experiments 1 to 16 used animal categories. Experiment 17 examined the centrality hypothesis with artifact categories. The results show centrality effects. Taken together, the Experiments suggest that centrality offers a domain-general constraint. Experiments 5, 8 to 11, and 17 left the properties that depended upon a candidate feature unspecified. A centrality effect was still obtained. The results suggest that centrality can operate under conditions of vagueness. The results are discussed in terms of theories of conceptual structure and models of category-based inference. A model to capture the present findings is also sketched.

ACKNOWLEDGMENTS

Many people made this thesis possible. First and foremost, I would like to thank my supervisor, Professor Rosemary Stevenson, for all her kind support and guidance. Big thanks go also across the Atlantic to Dr. S. A. Sloman and closer home to Professor D. E. Over for their stimulating discussions and critical comments throughout the preparation of this thesis. Thanks go also to Mr. Bob Metcalf for writing up the computer programs used in Experiments 12, 13, and 14.

I also would like to thank my family for their unconditional support (emotional and monetary). Many thanks go also to the Department of Psychology for awarding me a studentship. Thanks go also to the staff of the Psychology Department of the University of Durham for providing a rich academic and social environment - I enjoyed football matches the most.

Finally, I would like to thank my girlfriend, Barbara Minorini, for always being there for me and for proving to me that $1+1$ may exceed 2 in many different ways (carrying our baby being the best of them all!).

TABLE OF CONTENTS

ABSTRACT.....	2
ACKNOWLEDGMENTS	3
TABLE OF FIGURES.....	7
TABLE OF TABLES.....	7
DECLARATION	9
CHAPTER I: INTRODUCTION.....	10
1.1 THE 'PROBLEM' OF INDUCTION.....	12
1.2 CATEGORY-BASED ARGUMENTS	12
1.3 MAIN EFFECTS OF CATEGORY-BASED INFERENCE.....	14
1.3.1 <i>Category effects</i>	14
1.3.2 <i>Property effects</i>	16
1.3.3 <i>Property by category effects</i>	17
1.4 MODELS OF CATEGORY-BASED INFERENCE.....	19
1.4.1 <i>Category-based inference models for blank predicates</i>	19
1.4.2 <i>Category-based inference models for non-blank predicates</i>	23
1.4.3 <i>Summary of category-based inference models</i>	26
1.5 CONCEPTUAL CENTRALITY AS CENTRALITY IN 'INTUITIVE THEORIES'	27
1.6 CONCEPTUAL CENTRALITY AS MUTABILITY	30
1.7 MUTABILITY VERSUS VARIABILITY, DIAGNOSTICITY, AND SALIENCE	31
1.7.1 <i>Conceptual arguments</i>	31
1.7.2 <i>Empirical evidence</i>	33
1.8 MUTABILITY AS CENTRALITY IN A NETWORK OF PAIR-WISE DEPENDENCY RELATIONS	33
1.9 CONCEPTUAL CENTRALITY AND CATEGORY-BASED INFERENCE	35
1.9.1 <i>What we do not claim</i>	37
1.10 OBJECTIVE 1: ASSESSING THE CENTRALITY HYPOTHESIS	38
1.10.1 <i>Manipulations of conceptual centrality</i>	39
1.10.2 <i>Manipulations of shared structure</i>	42
1.11 OBJECTIVE 2: MUTABILITY VS. VARIABILITY	44
1.12 OBJECTIVE 3: INFERENTIAL PROCESS: DOMAIN GENERAL OR DOMAIN SPECIFIC?.....	45
1.13 OVERVIEW OF THE EXPERIMENTAL CHAPTERS	47
CHAPTER II: TWO MEASURES OF CONCEPTUAL CENTRALITY: IMMUTABILITY AND CENTRALITY IN A DEPENDENCY SPACE	48
2.1 MOTIVATION.....	50
2.1.1 <i>Categorization: Differential weighting of deep vs. surface features</i>	51
2.1.2 <i>Category-based inference: Projectible vs. nonprojectible properties</i>	51
2.2 BASIC EXPERIMENTAL DESIGN	52
2.3 EXPERIMENT 1: IMMUTABILITY	54
2.4 EXPERIMENT 2: MUTABILITY AND CENTRALITY IN A DEPENDENCY SPACE.....	58
2.5 GENERAL DISCUSSION	64
CHAPTER III: CONCEPTUAL CENTRALITY FROM A DEPENDENCY CHAIN	66
3.1 SHARED DEPENDENCY STRUCTURE	68
3.2 VAGUENESS AND DIRECTIONALITY	69
3.3 FEATURE CENTRALITY: CENTRALITY IN A CAUSAL OR A DEPENDENCY SPACE?.....	70
3.4 EXPERIMENT 3: CENTRALITY FROM A CAUSAL + EXPLANATION RELATION	71

3.5	EXPERIMENT 4: CENTRALITY FROM A CAUSAL RELATION.....	76
3.6	EXPERIMENT 5: CENTRALITY-IN-THE-ABSTRACT.....	78
3.7	EXPERIMENT 6: CENTRALITY FROM A SINGLE DEPENDENCY CHAIN	80
3.8	EXPERIMENT 7: REPLICATION OF EXPERIMENT 6.....	86
3.9	GENERAL DISCUSSION.....	89
3.9.1	<i>Centrality effect</i>	89
3.9.2	<i>Shared structure</i>	89
3.9.3	<i>Vagueness and directionality</i>	91
3.9.4	<i>Feature centrality: Centrality in a causal or a dependency space?</i>	91
3.9.5	<i>Blankety blank</i>	92
CHAPTER IV: CENTRALITY AND PROJECTIBILITY FROM THE "FEW/LOTS"		
DEFINITION.....		94
4.1	INTRODUCTION.....	95
4.1.1	<i>The "Few/Lots" definition and Sloman et al.'s (1998) notion of feature centrality</i>	95
4.1.2	<i>The "Few/Lots" definition helps control for necessity and sufficiency</i>	96
4.1.3	<i>The "Few/Lots" definition offers a different manipulation of shared structure</i>	97
4.1.4	<i>The "Few/Lots" definition provides a test for vagueness</i>	98
4.2	EXPERIMENT 8: FAMILIAR CATEGORIES, ABSTRACT PROPERTIES.....	99
4.3	EXPERIMENT 9: UNSPECIFIED CATEGORIES, ABSTRACT PROPERTIES .	103
4.4	EXPERIMENT 10: FAMILIAR CATEGORIES, CONCRETE PROPERTIES.....	106
4.5	EXPERIMENT 11: CENTRAL VERSUS LESS-CENTRAL VERSUS NO INFORMATION GIVEN PROPERTIES.....	112
4.6	GENERAL DISCUSSION.....	118
CHAPTER V: CENTRALITY VERSUS VARIABILITY: JUDGMENTS OF FREQUENCY AND PROJECTIBILITY.....		
		121
5.1	CONCEPTUAL CENTRALITY VERSUS VARIABILITY.....	122
5.2	A SUBTLE MANIPULATION OF FEATURE CENTRALITY.....	125
5.3	EXPERIMENT 12.....	126
5.4	EXPERIMENT 13.....	133
5.5	EXPERIMENT 14.....	138
5.6	GENERAL DISCUSSION.....	144
5.6.1	<i>Centrality versus frequency contributions on projectibility</i>	144
5.6.2	<i>Centrality information may bias the perception of frequency</i>	145
5.6.2	<i>Centrality and projectibility</i>	145
CHAPTER VI: CENTRALITY VERSUS VARIABILITY: INFORMATION SEARCH. 147		
6.1	CONCEPTUAL CENTRALITY VERSUS VARIABILITY: INFORMATION SEEKING.....	148
6.2	EXPERIMENT 15: CENTRALITY VS VARIABILITY IN INFORMATION-SEEKING.....	150
6.3	EXPERIMENT 16: A REPLICATION OF EXPERIMENT 15.....	153
6.4	GENERAL DISCUSSION.....	156
CHAPTER VII: EXTENDING THE CENTRALITY HYPOTHESIS TO ARTIFACTS.. 158		
7.1	NATURAL KINDS VERSUS ARTIFACTS.....	159
7.1.1	<i>Conceptual evidence and evaluation</i>	159
7.1.2	<i>Empirical evidence and evaluation: Type of property by type of kind interaction</i>	160
7.2	CONCEPTUAL CENTRALITY CAN ACCOUNT FOR CATEGORY BY PROPERTY INTERACTION.....	162
7.3	EXPERIMENT 17: IT'S ALL GREEK TO ME.....	163

7.4	GENERAL DISCUSSION.....	170
CHAPTER VIII: FINAL DISCUSSION.....		172
8.1	SUMMARY AND DISCUSSION OF FINDINGS.....	173
8.1.1	<i>Objective 1: All else being equal, the more central a feature for a concept the higher its projectibility to other concepts that share its structure.....</i>	173
8.1.2	<i>Objective 2: Conceptual centrality versus variability.....</i>	178
8.1.3	<i>Objective 3: Property induction: Domain specific or domain general ? ..</i>	178
8.1.4	<i>Objective 4: Vagueness.....</i>	179
8.2	RELATIONS OF CURRENT FINDINGS TO MODELS OF CATEGORY-BASED INFERENCE.....	179
8.2.1	<i>Modeling centrality effects.....</i>	180
8.2.2	<i>Toward a complete model of category-based inference.....</i>	185
8.2.3	<i>Relations between category-based models.....</i>	187
8.3	SUPERORDINATE STRUCTURE.....	189
8.3.1	<i>Hierarchies, judged similarity and property induction.....</i>	189
8.3.2	<i>Hierarchies, judged similarity and the effect of centrality on projectibility.....</i>	190
8.3.3	<i>Hierarchies, dependency structure, centrality and property induction.....</i>	191
8.4	SAMPLE SIZE AND MULTIPLE PREMISE ARGUMENTS.....	192
8.4.1	<i>Sample size.....</i>	192
8.4.2	<i>Multiple Premise Arguments.....</i>	193
8.5	IS FOLLOWING THE CENTRALITY HEURISTIC ADAPTIVE?.....	194
8.5.1	<i>Adaptiveness as maximizing real world success.....</i>	195
8.5.2	<i>Adaptiveness as minimizing psychological conflict.....</i>	199
8.6	ON CONCEPTUAL ASSUMPTIONS.....	200
8.6.1	<i>The multiple descriptions problem.....</i>	200
8.6.2	<i>Instability across tasks.....</i>	203
8.6.3	<i>Instability across development.....</i>	205
8.6.4	<i>Instability across cultures.....</i>	207
8.6.5	<i>Summary.....</i>	209
8.7	CONCEPTUAL CENTRALITY AND ESSENTIALISM.....	210
8.7.1	<i>Central features are not necessarily defining features.....</i>	211
8.7.2	<i>Central features are not necessarily deep features.....</i>	212
8.7.3	<i>Conceptual centrality depends on (specific or vague) knowledge.....</i>	212
8.7.4	<i>Conceptual centrality is domain-independent.....</i>	213
8.8	THE PROBLEM OF INDUCTION.....	213
APPENDIX.....		215
	<i>Chapter II Tables A.2.1 to A.2.9.....</i>	216
	<i>Chapter III Tables A.3.3 to A.3.10.....</i>	220
	<i>Chapter IV Tables A.4.1 to A.4.8.....</i>	227
	<i>Chapter V Tables A.5.1 to A.5.10.....</i>	229
	<i>Chapter VI Tables A.6.1 to A.6.5.....</i>	233
	<i>Chapter VII Tables A.7.1 and A.7.2.....</i>	235
	<i>Chapter VIII Table A.8.1 and A.8.2.....</i>	236
	<i>Contents of the accompanying floppy disk (It contains Excel files of all the data).....</i>	237
BIBLIOGRAPHY.....		238

TABLE OF FIGURES

Figure 1. 1	Illustration of the feature-coverage model for the argument "Robins have an ulnar artery; therefore falcons have an ulnar artery".....	21
Figure 7. 1	A simplified illustration of the view that dependencies point inwards for natural kinds but outwards for artifacts. Dots represent properties. (adapted from Keil, 1995).	171
Figure 8. 1	Simplified illustration of a generic dependency structure, and a special one reflecting essentialism. Dots stand for properties.....	211

TABLE OF TABLES

Table 2. 1	The full set of materials of Experiment 1. The left panels of the table present the premise-conclusion categories and the right panels the properties used.	55
Table 2. 2	Mean (SE) centrality estimates for the central and less central items in each task order for Experiment 1.....	57
Table 2. 3	Mean (SE) centrality estimates for the central and less central items for each centrality measure for Experiment 2.	61
Table 2. 4	Mean conditional (unconditional) likelihood estimates as a function of Centrality and Centrality Estimation task.	62
Table 3. 1	Design and Materials of Experiments 3 and 4. Experimental = causal+explanation condition (Exp. 3), causal condition (Exp. 4).....	72
Table 3. 2	Sample items from each Shared Structure by Relation condition of Experiments 3 and 4.	73
Table 3. 3	Mean proportion choice of the property designed as the cause in the causal+explanation condition.	74
Table 3. 4	Mean (SE) proportions of the property designed as the cause in the Causal condition.	77
Table 3. 5	Animal pairs of Experiment 10 along with the mean (SE) similarity ratings....	81
Table 3. 6	Sample items from each Shared Structure by Centrality condition of Experiment 6. Only one name order and one premise triple by property type assignment is shown.	82
Table 3. 7	The 18 premise-conclusion categories of Experiments 6, 7, 10, 11 (first 3 triples), 15, and 16.....	83
Table 3. 8	Design of the 3 lists of materials of Experiments 6, 7, and 10.	83
Table 3. 9	Properties used in Experiments 6 and 7.....	84
Table 3. 10	Mean (SE) inductive strength (mean likelihood estimates) for the Shared structure by Centrality conditions.	85
Table 3. 11	Mean (SE) inductive strength (mean likelihood estimates) for the Shared structure by Centrality conditions. Super=Superordinate.	88
Table 4. 1	Sample items from each Shared Structure condition for Experiment 8. ...	100
Table 4. 2	Animal pairs for Experiment 8, List 1.....	101
Table 4. 3	Mean (SE) percentage choice of the central property.	102
Table 4. 4	Mean (SE) inductive strength estimates for Experiment 9.....	105
Table 4. 5	Sample items from each Shared Structure by Centrality condition of Experiment 10.	107
Table 4. 6	Design of the 6 lists of materials of Experiment 10.	108
Table 4. 7	Mean (SE) inductive strength estimates as a function of Centrality and Shared structure for Experiment 10.....	109

Table 4. 8	Correlation statistics for items between estimates of shared structure and estimates of inductive strength for central properties, less central properties, and their difference.	111
Table 4. 9	Sample items from each Shared Structure by Centrality condition of Experiment 11.	115
Table 4. 10	Mean (SE) inductive strength as a function of Similarity and Centrality for Experiment 11.	115
Table 5. 1	Mean (SE) inductive strength estimates for the novel properties of Experiment 12.	129
Table 5. 2	Actual frequency estimates and mean (SE) frequency estimates for each task order for Experiment 12.	131
Table 5. 3	Mean (SE) inductive strength estimates for the familiar properties of Experiment 12.	132
Table 5. 4	Actual frequency estimates and mean (SE) judged frequency estimates for the familiar properties of Experiment 12.	132
Table 5. 5	Mean inductive strength estimates for Experiment 13.	135
Table 5. 6	Actual frequency estimates and mean (SE) frequency estimates for each task order for Experiment 13.	136
Table 5. 7	Mean (SE) inductive strength estimates for the familiar properties of Experiment 13.	137
Table 5. 8	Actual and mean (SE) judged frequencies for the familiar properties of Experiment 13.	137
Table 5. 9	Mean (SE) percent inductive strength estimates for the novel properties averaged across frequency conditions of Experiment 14.	140
Table 5. 10	Actual and mean (SE) frequency estimates for each task order for Experiment 14.	141
Table 5. 11	Mean (SE) inductive strength estimates for the familiar properties of Experiment 14.	142
Table 5. 12	Frequencies and mean (SE) frequency estimates for the familiar properties for Experiment 14.	142
Table 6. 1	Mean choice of Centrality information in each similarity condition. Chance level = 3.00.	151
Table 6. 2	Sample items from Experiment 15.	154
Table 6. 3	Mean choice of centrality information in each shared structure condition. Chance level = 3.00.	155
Table 7. 1	Premise triples along with the target categories for the items of Experiment 17.	166
Table 7. 2	Mean (SE) functional similarity and surface similarity ratings for the 3 similarity conditions. The higher the rating, the higher the judged similarity.	167
Table 7. 3	Mean (SE) inductive strength estimates for Centrality by Similarity conditions.	168

DECLARATION

The research contained in this thesis was carried out by the author between October 1996 and December 1999 whilst a postgraduate in the Department of Psychology at the University of Durham. None of the work contained in this thesis has been submitted in candidate for any other degree.

STATEMENT OF COPYRIGHT

The copyright of this thesis remains with the author. No quotation from it should be published without his prior written consent and information derived from it should be acknowledged.

WORK FOR THIS THESIS HAS RESULTED IN THE FOLLOWING PUBLICATIONS:

Hadjichristidis, C., Sloman, S. A., Stevenson, R. J., and Over, D. E. Centrality and property induction (1999). *Proceedings of the Twenty first Annual Conference of the Cognitive Science Society*. Vancouver, British Columbia, 795.

Hadjichristidis, C., Sloman, S. A., Stevenson, R. J., and Over, D. E. Centrality and property generalization. *Proceedings of the European Conference on Cognitive Science*. Italy, Siena, 185-190.

CHAPTER I:
INTRODUCTION

One of our fundamental capacities is using old beliefs to generate new ones. Whenever we do so, we have drawn an inference. Much everyday reasoning involves drawing inferences to reach conclusions about which we are uncertain. This type of inference is known as *induction* (for a good introduction on inductive inference see Skyrms, 1975). Induction is an umbrella term for "all inferential processes that expand knowledge in the face of uncertainty" (Holland, Holyoak, Nisbett, & Thagard, 1986, p.1). The present studies investigate *category-based induction* that involves generalizing properties among categories. We address why some properties are more generalizable than others and why the generalizability of a property (or of a type of properties) depends on the specific premise and conclusion categories used. We propose that property induction is constrained by the immutability of a feature for a concept, that is by the extent to which a feature is mentally transformable while retaining the concept's coherence. Under the supposition that concepts consist of features embedded in networks of asymmetric dependency relations, we assume that immutability can be surrogated by centrality in a concept's feature space. That is, we take a feature to be central (immutable) to the extent that other (central) features depend on it (cf. Sloman, Love, & Ahn, 1998). We investigate the hypothesis that the more central a feature in a category's representation, the higher its projectibility to other concepts that share its structure.

The discussion proceeds as follows. We unpack the problem of induction that we address, and present the terminology of category-based inference. Then we highlight some robust effects of category-based induction. A look at some formal models of category-based inference follows, where we discuss their assumptions, successes and limitations. This discussion lays the background, experimental framework, and motivation for our proposal. Subsequently, we present our proposal of conceptual centrality as an inductive constraint in detail - we clarify the representational assumptions and the way centrality is captured from those assumptions, we derive empirically testable hypotheses, and we motivate our research by conceptual and empirical evidence. Finally, we present the break down of studies into experimental chapters, and the hypotheses that each experimental chapter addresses.

1.1 THE 'PROBLEM' OF INDUCTION

The current research is an attempt towards solving the problem of induction. To be clear, our research is not about solving the traditional problem of induction, i.e. the problem of whether induction can be logically justified (e.g. Hume, 1748/1988; Bonjour, 1992). Our starting point is that it is surely rational or reasonable to engage in inductive reasoning because otherwise we would not be able to learn from experience or to have expectations about the future - all of our evidence would be about the past (see Harman, 1995). The problem of induction that we address is this: Why are we more willing to generalize certain (types of) properties among categories rather than others? Why and how does the generalizability of (types of) properties depend on the (kinds of) categories involved? In the words of Mill (1843/1974): "Why is a single instance, in some cases, sufficient for a complete induction, while in others myriads of instances, without a single exception known or presumed, go such a little way toward establishing a universal proposal?" (p. 314, cited in Nisbett, Jepson, Krantz, & Kunda, 1983, p. 342). In Goodman's (1955) terms, why are some properties more projectible than others? Why are we more confident, for instance, that a newly discovered mammal will have a *heart* rather than *fur*? Similarly, why are we more willing to infer that a new refrigerator will *freeze food* rather than it will be *white*?

--BACKGROUND ON CATEGORY-BASED INFERENCE--

1.2 CATEGORY-BASED ARGUMENTS

Arguments are statements with a set of premises followed by a conclusion. *Categorical arguments* are arguments of the general form *Members of category A have property X, therefore members of category B have property X¹*, where property X remains constant across the premise and conclusion categories, and A and B are (psychologically) simple categories like robin, tomato, and bicycle. For instance, *Eagles have an ulnar artery*;

¹ Henceforth arguments and predicates are italicized.

therefore falcons have an ulnar artery, is a categorical argument. Categorical arguments are frequently depicted vertically as in the example below,

Eagles have an ulnar artery

Falcons have an ulnar artery

The statement above the dotted line is the premise of the argument and is assumed to be true, and the statement below the dotted line is the argument's conclusion.

The predominant measure of the strength of a categorical argument is *argument strength* or *psychological strength*. Argument strength refers to the extent to which belief in the premises of an argument leads to belief in its conclusion. Mere conviction inspired by the conclusion of an argument (independent of its premise) is not sufficient for psychological strength. To make the point clear, consider the following argument:

Doves are white

Refrigerators are white

The conclusion of this argument is credible; most refrigerators are in fact white. However, this argument is psychologically weak since the belief in the argument's conclusion was not transmitted from belief in its premise. (For a detailed discussion of argument strength see Osherson, Smith, & Shafir, 1986).

Two main aspects can be distinguished in a categorical argument: its premise and conclusion categories and the candidate property. In vertical depictions, these parts refer respectively to the left- and the right-hand sides of a categorical argument. The relation between an argument's premise and conclusion categories, as well as aspects of the candidate property both seem to contribute (though not necessarily independently) to its strength. Below we highlight the main findings of category-based induction in 3 sections reflecting this distinction. Our discussion is limited to specific single-premise arguments since these are the

arguments the present research investigated. (For an up-to-date thorough review on inductive inference see Heit, in press).

1.3 MAIN EFFECTS OF CATEGORY-BASED INFERENCE

1.3.1 Category effects

Some researchers studied category effects on induction by using predicates about which their participants had few beliefs such as *have the neurotransmitter dihedron, secrete uric acid crystals, or have a communicable disease* (e.g. Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Rips, 1975; Sloman, 1993). The idea behind the use of such predicates (named *blank*) is that they allow us to study belief transmission between categories while keeping background knowledge effects at a minimum². Under this assumption arguments using blank predicates have been depicted in the format “premise category/conclusion category” suppressing the blank predicate used. The argument, *Falcons have an ulnar artery; therefore eagles have an ulnar artery*, for example, has been depicted as “Falcons/Eagles”. Properties and therefore property effects are not considered. Argument strength is captured by the conditional probability of the conclusion given the premise.

Studies using blank predicates repeatedly reveal that argument strength is an increasing function of three aspects of the argument's categories:

1. The similarity between the premise and conclusion category
2. The typicality of the premise category
3. The specificity (homogeneity) of the conclusion category

² Blank predicates have also the advantage of retaining a task's meaningfulness. That is, it is meaningful to ask participants for their confidence in projecting blank predicates since there must be a correct answer to such queries; e.g. eagles either do or do not have an ulnar artery. Asking participants about abstract predicates, such as *has substance z*, though it would (presumably) prevent participants from bringing background knowledge to bear in the task, it would increase the task's vagueness - it is unclear whether there is a right answer to such a question.

Premise-conclusion similarity. All else being equal, the more similar a conclusion category to the premise category the stronger the inference. This is known as the *premise-conclusion similarity* phenomenon and was first discovered by Rips (1975). Rips asked adult participants to imagine that on an island all members of a particular species of animals (the premise category) had a new type of contagious disease. Participants were then asked to estimate the proportion of various other species of animals (the target categories) that also had the disease. Rips found that the greater the similarity between the premise and the target categories (measured as a distance in a multidimensional scale solution) the stronger the inference. For instance, inferences from rabbits to dogs were stronger than inferences from rabbits to bears. The premise-conclusion similarity effect is very robust. It has been demonstrated with adults (e.g. Osherson et al., 1990; Osherson, Stern, Wilkie, Stob, and Smith, 1991; Sloman, 1993), young children (e.g. Carey, 1985; Gelman, 1988), infants (e.g. Baldwin, Markman, & Melartin, 1993; Mandler and McDonough, 1996), and highly diverse populations such as North American students and Itzaj Mayans (e.g. Lopez, Atran, Coley, Medin, & Smith, 1997).

Premise typicality. All else being equal, the more typical a premise category to its superordinate category, the stronger the inference. This is known as the *premise-typicality* phenomenon. Rips (1975), for instance, found that arguments having bluejay (a typical American bird) as a premise category were rated as stronger than arguments having goose as a premise category. The premise typicality effect is also very robust. (For further studies with adults see Osherson et al, 1990; Osherson et al., 1991; for studies with young children see Carey, 1985; Lopez, Gelman, Gutheil, and Smith, 1992; for cross-cultural studies see Lopez et al., 1997).

The effects of premise typicality are not reducible to the effects of premise-conclusion similarity (at least for models that view similarity as a symmetric relation). The best evidence for this comes perhaps from the *premise-conclusion asymmetry* phenomenon. Inferences from typical to atypical categories (e.g. an inference from robins to chicken) are generally judged stronger than inferences from atypical to typical categories (e.g. an inference

from chickens to robins) even though the premise-conclusion similarity remains fixed (see e.g. Carey, 1985; Osherson et al., 1990).

Conclusion specificity. Finally, all else being equal, the more specific - the lower down in the taxonomy - a conclusion category (e.g. Belgian shepherd is a more specific category than dog which is more specific than mammal) the stronger the inference. This is known as the *conclusion specificity* phenomenon (Osherson et al., 1990). Osherson et al. (1990) found, for instance, that an argument having bird as a conclusion category was rated stronger than an argument having animal as a conclusion category. (For related evidence that the scope of a conclusion category constrains induction see McDonald, Samuels, & Rispoli, 1996). To the extent that more specific categories are more homogeneous, the more general conclusion might be that the more homogeneous a conclusion category the stronger an inference (see Gelman, 1988; Heit, in press; Nisbett et al, 1983).

To summarize, categories do matter a lot in category-based inference. Specifically, the similarity between the premise and conclusion categories, the typicality of the premise category, and the homogeneity of the conclusion category all promote strong arguments

1.3.2 Property effects

Specifics of the candidate properties also play a critical role on judged argument strength. Some properties are more projectible than others (for evidence with adults see Nisbett et al, 1983; for evidence with children see Gelman, 1988, and Gutheil & Gelman, 1997; for evidence with infants see Mandler & McDonough, 1998). Nisbett et al. (1983), for instance, asked participants to imagine that they were explorers on a remote island where they observed a sample of instances (animals, objects, or people) all of which had a property. Their task was to guess the percentage of other instances of the same type having that property. For instance, participants were informed that 3 individuals of a tribe were observed all of who were *obese* and had *brown skin*. Based on this fact they were asked to estimate the percentage of the individuals of the tribe that shared each characteristic. Participants judged that a higher percentage of individuals would have brown skin than they would be obese. Nisbett et al.

claimed that this was because the former property was perceived as more homogeneous than the latter. Their claim was supported by participants' reports.

In summary, on the top of categories, candidate properties also matter a great deal. But what decides which properties are projectible and which are not? To say that some properties are more projectible because they are believed to be more homogeneous is to say no more than humans are able to reason statistically - this in fact was the point that Nisbett et al. (1983) wanted to make with their study.

1.3.3 Property by category effects

Projectibility is not an attribute that features have in vacuum. Rather the effect of a property of an argument critically depends on the argument's categories. For instance, *color* may be fairly projectible between different types of polar animals, but not between different types of animals. Generally, evidence suggests that the effect of a (kind of) property on projectibility depends on the extent to which an argument's categories are similar in terms of the relations and attributes that are meaningfully connected to the candidate property. (For evidence with adults see Heit & Rubinstein, 1992; Ross & Murphy, 1999; Sloman, 1994, 1997; Smith, Shafir, & Osherson, 1993; for evidence with children see Gelman & Markman, 1986; Kalish & Gelman, 1992).

Heit and Rubinstein (1994), for instance, found that for the anatomical property *has a liver with two chambers* inferences were stronger from hawks to chickens than from tigers to chickens, whereas for the behavioral property *prefers to feed at night* the order of the preference was reversed (cited in Heit, 1997). They claimed that there are (at least) two distinct types of similarity: anatomical and behavioral. An inference is strong to the extent that the type of the candidate property matches the type of similarity between the premise and conclusion categories. Notice that to the extent that the properties Heit and Rubinstein considered were blank, their results challenge the existence of blank predicates; perhaps all predicates enter the reasoning process.

But what does it mean to be meaningfully related to a candidate feature? What types of relations and properties matter? Lassaline (1996), using both abstract and concrete animal categories, showed that a feature is projectible to the extent that it is causally related to a property that the target shares. The strength of the argument:

Animal A has loose tensor tympani and frequent migraines
Animal B has loose tensor tympani

Animal B has frequent migraines

should increase by adding the premise *For animal A loose tensor tympani cause frequent migraines*. (The properties in this example are my own). Notice that without the added premise Animal A and Animal B share a property, whereas with the added premise they share a causal antecedent. Lassaline's studies provided evidence that argument strength is constrained by the principle of systematicity (see Gentner, 1983, 1989). According to this principle, a match of a connected system of relations is preferred to a match of an equal number of unconnected relations (this point will be elaborated when we will present structural-mapping models of category-based inference).

Keil and his colleagues (in preparation, cited in Keil, 1995) have shown an abstract type of property by type of category interaction in category-based inference with children, an interaction that can not be attributed to (specific) causal knowledge. A group of children were presented, for instance, with an ambiguous picture labeled as 'rock' while another group with the same picture labeled as 'frog'. Participants were shown two other pictures and had to choose the one that belonged to the same category as the target object. The results showed that when the property was labeled as 'frog', shape rather than color and surface markings were important. When the same entity was labeled as 'rock', color and surface markings became more important. Keil (1995) presented studies that showed a similar property by category interaction while using pictures vaguely described as either novel animals, or novel machines. Once again it seems that the projectibility of a property depends on the respect to which the premise and conclusion categories are similar. Critically because the objects that

Keil and his colleagues used were by definition novel, the knowledge underlying children's judgments should be in the form of abstract pre-theoretical biases rather than of concrete causal knowledge.

To summarize, properties (perhaps even blank properties) enter the reasoning process by potentiating some aspects of the premise category. Arguments are strong to the extent that the conclusion category shares those aspects. The general case seems therefore to be that the effect of properties on projectibility is intertwined with the effect of categories. A critical question is what types of properties and relations does a candidate feature potentiate. Lassaline (1996) and Wu and Gentner (1998) claim it is those properties that are causally related to the candidate feature. Sloman (1994, 1997) and Smith et al. (1993) assume (at least for explainable predicates) it is properties that are relevant in explaining the predicate. Interestingly, Keil and his colleagues demonstrated that some property effects cannot be explained in terms of specific causal knowledge but rather in terms of vague pre-theoretical biases.

1.4 MODELS OF CATEGORY-BASED INFERENCE

Having discussed the three main findings associated with single-premise categorical arguments, we turn to various models of category-based inference and briefly discuss the extent to which each addresses these basic effects.

1.4.1 Category-based inference models for blank predicates

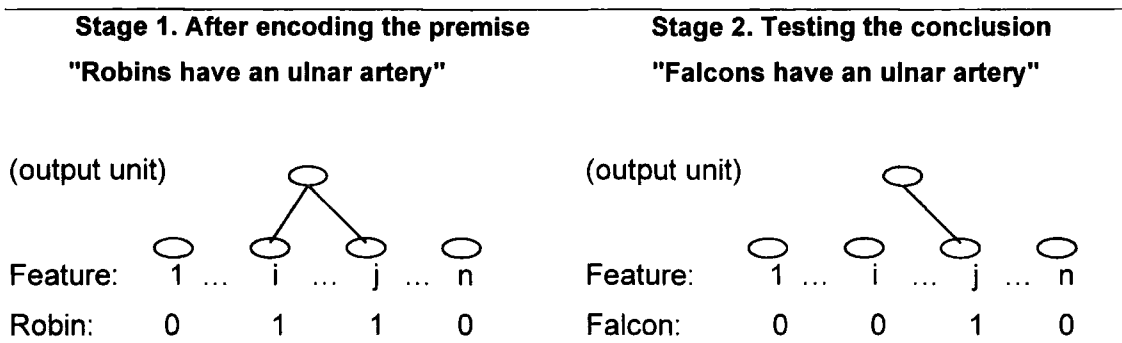
The similarity-coverage model. Osherson et al. (1990) advanced the *similarity-coverage model* (SCM) as an attempt to capture category-based induction with blank predicates. The model assumes that categories are represented by exemplars and that they are structured in stable hierarchical representations. The model predicts that argument strength increases with (i) the degree to which the premise and conclusion categories are similar (the similarity component), and (ii) the degree to which the premise categories are similar to the

lowest level category that includes both the base and conclusion categories (the coverage component). In the special case of single-premise arguments premise coverage reduces to the typicality of the premise category to the least inclusive superordinate category.

The model can capture all three main category effects. The similarity component directly captures premise-conclusion similarity effects. The coverage component directly accounts for premise typicality effects. The model can also capture some conclusion-specificity effects to the extent that the more specific a conclusion category the better the premise covers it (e.g. robins cover better the category birds than the more inclusive category animals). In fact, the model can successfully account for an impressive array of a dozen or so phenomena associated with blank biological predicates.

The feature-coverage model. As an alternative to Osherson et al.'s similarity-coverage model, Sloman (1993) advanced the *feature-coverage model* (FCM). The FCM diverges from the similarity-coverage model in that it does not assume a hierarchical structure, but instead it represents categories at all levels (i.e. basic (e.g. Elephant), subordinate (e.g. Indian Elephant), and superordinate (e.g. Mammal)) as vectors of values over a set of features. The model consists of a network of n features (input units), which are used to represent categories, and a single output unit, which is used to represent the candidate property (see Figure 1.1). The model assumes a two-stage process. During the first stage, the premise category is encoded as a vector of weights by connecting all the features that represent the premise category to the candidate property. In the second stage, the conclusion is tested by examining the activation of the candidate property (output unit) upon the presentation of the features that represent the conclusion category. Figure 1.1 illustrates the process. Roughly, an argument is predicted as strong to the extent that the conclusion category shares features with the premise category and has a few distinctive features of its own. The feature-coverage model takes advantage of the automatic generalization of distributed representations to project a property between categories on the basis that the features of the premise categories cover those of the conclusion category.

Figure 1. 1 Illustration of the feature-coverage model for the argument "Robins have an ulnar artery; therefore falcons have an ulnar artery".



Note. A feature value of 0 means that the category does not have the property, a value of 1 means that it does.

The FCM can also account for the three main category effects. It readily captures premise-conclusion similarity effects since the activation of the predicate upon the presentation of the conclusion category is proportional to the premise-conclusion feature overlap. The model can also account for premise typicality effects to the extent that atypical conclusion categories (e.g. ostrich) have more idiosyncratic features than typical conclusion categories (e.g. robin). Finally, the FCM could account for conclusion-specificity effects to the extent that the more specific a conclusion category the more properties it shares with the premise. In fact, the FCM can account for roughly the same array of phenomena associated with blank predicates as the SCM.

Sloman (1998) provided further empirical support for his model by showing that, at least under certain circumstances, participants appear to ignore the hierarchical structure of categories. For example, many participants did not assign a maximum strength for inferences from birds to robins or from birds to ostriches, but rather their estimates were proportional to the premise-conclusion similarity. This result directly supports the feature-coverage model and challenges the similarity-coverage model's assumption of stable hierarchies.

Hampton (1982) provides further evidence against the idea of pre-stored hierarchies by demonstrating intransitivities. He found, for instance, that people are willing to accept that *car seats* are kinds of *chairs* and that *chairs* are kinds of *furniture*, but denied that *car seats* are kinds of *furniture*. If concepts were hierarchically organized in memory, the claim goes,

and furniture is a superordinate of chairs which is a superordinate of car seats, then it should automatically follow that car seats are kinds of furniture. (For further evidence of intransitivity see Kempton, 1978; Randall, 1976).

Limitations. Although the Osherson et al. (1990) SCM and the Sloman (1993) FCM models can capture an impressive array of phenomena, their scope is limited. They cannot (nor do they claim to) capture property or property by category effects. To capture property by category effects a model should at least allow flexibility in the comparison between the premise and conclusion categories. To capture the findings of Lassaline (1996), for example, a model should represent somehow causal inter-property relations and offer a mechanism to compare such structured representations. However, neither model assumes relational structure nor is flexible in comparing representations.

The need to capture relational structure is pressing because conceptual arguments strongly suggest that category representations cannot simply involve a summation of features. A summation of features such as *barks*, *has a tail*, and *has two eyes*, for instance, would fail to capture dog since these properties in a weird arrangement, say putting the eyes in the tail, would not sum up to a dog but perhaps to a Picasso painting (Hahn & Chater, 1997, see also Murphy & Medin, 1985). Taking the argument from a different angle, to the extent that category-based inference is determined by similarity (like the SCM and the FCM assume) and similarity is influenced by inter-property relations, then argument strength should also be influenced by such relations. Recent experimental evidence suggests that relations do influence similarity judgments (Goldstone, Medin, & Gentner 1991; Goldstone, 1994b). Hence, relations should also influence category-based inference.

The scope of the SCM and FCM models seems even more limited in the face of Heit and Rubinstein's (1994) findings of property by category effects for properties that (presumably) fall within the models' pre-described domain (i.e., the domain of blank

predicates). Heit and Rubinstein's findings are damaging to both models to the extent that the properties they considered were blank³.

To summarize, without doubt the similarity-coverage model and the feature-coverage model have increased our understanding of category-based inference. However, the scope of these models is limited. It is uncertain how these models can scale up to deal with non-blank predicates.

1.4.2 Category-based inference models for non-blank predicates

Structural-mapping models. In contrast to models that view concepts as composed of exemplars or unstructured lists of attributes, structural-mapping models view category representations as composed of features embedded in hierarchical structures of relations (e.g. Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983, 1989; Goldstone, 1994a; Goldstone & Medin, 1994 a, 1994 b; Markman & Gentner, 1993 a, 1993 b). Object comparison is seen as a process of *structural alignment* between representations; a process that puts the representations of the premise (commonly known as the *base* in this literature⁴) and

³ According to Osherson et al. (1990) blank predicates "...involve predicates about which subjects in our experiments have few beliefs." (p. 186). According to this definition, the blankness of a predicate for a given individual depends on whether that individual is familiar with the predicate. Some of the predicates that Heit and Rubinstein consider, e.g., *prefers to feed at night* seem to fall short of the mark. Yet it is still unclear how one could independently measure property blankness. One possibility is to ask for the unconditional probability of the premise and conclusion statements; e.g., How likely is it that eagles have an ulnar artery? Even if such probability estimates are gathered, how is one to interpret them? Does no knowledge equate with a probability of 0, .5, or any probability whatsoever?

Sloman and Wisniewski (1992) provide a different definition of blankness: a predicate is blank as long as people cannot explain the relation between predicate and category. Notice that this definition diverges from Osherson et al.'s in that it emphasizes the need to consider a property as it relates to categories rather than in isolation. In the argument *Collies hate salted peanuts, therefore Siamese cats hate salted peanuts*, for instance, *hate salted peanuts* is a familiar but unexplainable predicate and hence according to their definition blank. Sloman and Wisniewski provided empirical evidence that the domain of the *feature-coverage* model (and presumably the domain of the *similarity-coverage* model as well) can be extended as to include familiar but unexplainable predicates (for inductive inference about explainable predicates see Sloman, 1994, 1997).

In a circular sense, the properties that Heit and Rubinstein studied are not blank according to the Sloman and Wisniewski definition since they produced different inductive strength orderings for arguments that shared premise and conclusion categories. To avoid such circularity, one needs to devise a test for blankness independent of the process studied. It is not clear how such a test can be devised.

⁴ Structural-mapping models are most associated with literature on analogy. Such models, however, readily extend to cases of literal similarity. Literal similarity and analogy are seen as two types (continua) of similarity the first one involving lots of relational and attribute matches, while the latter lots of relational but few attribute matches (see e.g. Gentner, 1989).

conclusion (commonly known as the *target* in the same literature)⁵ categories into a correspondence (alignment). Inference is seen as a direct consequence of this structural-mapping process.

The structural-alignment process is carried out by the *structural-mapping engine* (SME) which determines the maximally *structurally consistent alignment* between two representations. A structurally-consistent alignment is one that obeys (1) a *one-to-one mapping* (each element of the base maps to at most one element of the target object), and (2) *parallel connectivity* (if one element of the base is matched with an element of the target, then their arguments must correspond as well). The maximally structurally consistent mapping is determined by the principle of systematicity. According to this principle there is a built-in bias to show a preference to match common rather than isolated systems of relations (Gentner, 1983, 1989). A match of a connected system of relations is preferred to a match of an equal number of unconnected relations.

At the end of the comparison process the system is left with a maximally structurally consistent mapping between the representations. Properties that are connected to the resulting system of relations that are present in the base but absent in the target constitute potential inferences (see Clement & Gentner, 1991; Lassaline, 1996, Wu & Gentner, 1998). Following the principle of systematicity, the goodness of a candidate inference increases in proportion to the order of the relation that it is embedded in - the higher the order the stronger the inference.

Structural mapping models can directly account for the results of Lassaline (1996) and Wu and Gentner (1998). Moreover such models have been also said to account (Gentner & Medina; 1998; Wu & Gentner, 1998) for the various property by category effects such as those reported by Heit and Rubinstein (1994). When the candidate property is behavioral (or anatomical), the claim goes, participants search for meaningful properties in the base that are causally related to this property (presumably behavioral or anatomical). An inference is strong to the extent that the target shares these properties with the base. It is less obvious how the

⁵ Henceforth, the terms premise and base, as well as the terms conclusion and target will be used interchangeably.

structural mapping models can account for category effects. Assuming that in the absence of causal relations the degree of matching between the premise and conclusion categories would be proportional to the extent that their features overlap, the structural mapping models should be able to account for premise-conclusion similarity effects and perhaps, because of that, also for some premise-typicality and conclusion-specificity effects. Finally, the structural mapping models cannot readily capture effects of abstract pre-theoretical biases such as those reported by Keil and his colleagues (e.g. Keil, 1995).

The gap model. The gap model (Smith, Shafir, & Osherson, 1993) was an attempt to extend reasoning models to non-blank predicates such as *can bite through barbed wire*. The model assumes that non-blank predicates enter the reasoning process by potentiating the dimensions of the premise category that are meaningfully related to the predicate. The theorists assumed, for instance, that an argument with the premise *Poodles can bite through barbed wire* would potentiate the dimensions *strength* and *ferocity*. Smith et al. also assume that predicates are associated with some criterion value on those dimensions. Non-blank predicates invite an examination of the plausibility of the premise. If the values on the potentiated dimensions of the premise category are lower than the criterion value for satisfying the predicate, the values of the predicate would be lowered; the bigger the gap, the bigger the downward adjustment (e.g. the premise *Poodles can bite through barbed wire* would invite more downward adjustment of the predicate values than the premise *Dobermans can bite through barbed wire*). Broadly, the closer the values on the respective dimensions of the conclusion category are to satisfy the adjusted criterion values of the predicate, and the more similar the premise and conclusion categories are, the stronger the argument.

The gap model can account for some property by category effects. For instance, it can readily account for why participants find the argument *Poodles can bite through barbed wire; therefore, German shepherds can bite through barbed wire* stronger than the argument *Dobermans can bite through barbed wire; therefore, German shepherds can bite through barbed wire* (see Smith et al., 1993). Because it incorporates a similarity component, the gap model can also account for premise-conclusion similarity effects. Further, because similarity

is related to typicality - e.g. German shepherd is both more similar and more typical to dog than chow-chow is - the gap model with its similarity component can perhaps also capture some premise-typicality effects. The main limitation of the gap model is that even though it critically depends on a feature potentiation process, it presupposes rather than states how this process takes place.

1.4.3 Summary of category-based inference models

Both Osherson et al.'s (1990) similarity-coverage model and Sloman's (1993) feature-coverage model are successful in that with a minimal theoretical edifice they can account for a wide array of phenomena associated with blank biological predicates. These models, however, cannot account for relational effects because they do not assume relational structure. Irrespective of whether unfamiliar properties embedded in causal relations (such as those studied by Lassaline, 1996) are blank, it is transparently clear that a complete model of property induction, to be viable, must account for such properties' effects. In this respect, models viewing inference as resulting from a structural-mapping process have an advantage. Such models offer principles of how the goodness of an inference is constrained by a mapping process that heavily weights relational matches, thus directly addressing some property by category effects. The gap model (Smith et al. 1993) can also capture some effects associated with non-blank predicates but most of its explanatory burden rests in the potentiation process, a process that the model assumes rather than explains. Finally, both structural mapping models and the gap model can, to a certain extent, account for the three main category effects because both view argument strength as an increasing function of shared features.

Is there any underlying commonality to all these models? At a broad level, these models rely on similar constructs to explain category-based inference. The Osherson et al. (1990) SCM, the Sloman (1993) FCM and the Smith et al. (1993) gap model all put a heavy explanatory burden on similarity. The gap and the structural mapping models can be said to maximize coherence. A central principle of the structural mapping models is that of systematicity. The principle of systematicity can be directly translated as a principle of

maximizing coherence. The gap model can also be said to maximize coherence. The notion underlying the gap model is that a conclusion category is tested on the closest possible world where the premise category satisfies the predicate. To the extent that the closest possible world is one that satisfies the premise while retaining the relations between the premise and conclusion categories that hold on the actual world, the gap model is maximizing coherence. The present account also views category-based inference in terms of coherence maximization - albeit a sort of coherence that is not discussed by the reviewed models.

--BACKGROUND ON CONCEPTUAL CENTRALITY--

1.5 CONCEPTUAL CENTRALITY AS CENTRALITY IN 'INTUITIVE THEORIES'

None of the category-based models reviewed provides constraints on what properties get tabulated and ultimately represented. Constraints on what properties get tabulated and represented are critical because they point to constraints of inductive inference as well - aspects of category representations should influence concept use. The aim of the present section is not to criticize the models reviewed - none of them purports to completely describe category-based inference. Further, the problem of feature selection is very hard (see Chater & Oaksford, 1993; Oaksford & Chater, 1991). Rather the discussion aims to motivate our hypothesis; i.e., that feature centrality influences property generalization.

None of the category-based inference models considered addresses what features get tabulated and represented in the first place. It is left up to the researchers to feed their models with the representations of the premise and conclusion categories. Otherwise stated, the models presuppose rather than explain the content of mental representations (Murphy & Medin, 1985). The researchers are left to rely on some independent measure to capture representational content. One such measure, the measure that both Osherson et al. (1991) and Sloman (1993) used, is *consensual validation* (see Rosch and Mervis, 1975). According to this, representations are taken to consist of properties that many people list for a particular

category. A problem with consensual validation is that it may produce different features depending on the level of abstractness of a category. People may list, for example, *means of transport* for vehicle but not for Honda Civic, though they know that this feature applies also to the former category. B. Tversky and Hemenway (1986) argued that the features that people list critically depend on their assumptions about the relevant *contrast set* (the set of categories that participants implicitly contrast a given category against) and the desired level of specificity of the category. Returning to the example, the feature *means of transport* is important for differentiating vehicles from categories such as furniture (given that furniture is in the relevant contrast set of vehicle), but not for differentiating Honda Civic from, say, Honda Legend. Put in other words, consensual validation downplays people's amount of knowledge about categories.

What factors constrain feature representation? Some of the constraints are provided by the perceptual system itself (see Ullman, 1979). To take an extreme example, since we cannot perceive quark particles (with the naked eye) there is no way that we can tabulate them. More interestingly, property tabulation, and therefore property representation as well, is biased by our theories, our intuitions and ideas about how the world works. There is an infinite number of co-occurrences that we could potentially tabulate, such as a relation between people's toe length and their favorite beverage. Our theories though predispose us to expect that no relation is to be found in such cases (see Heit, 1997), thus preventing us from tabulating and hence representing such co-occurrences.

Empirical evidence abounds that suggests lay theories influence how we perceive and interpret properties. A striking example is provided by Wisniewski and Medin (1991, 1994). Participants were given children's drawings to observe and were told that some of the drawings were done by city children while others by farm children. The experimenters randomly assigned these labels to a set of drawings and in most cases the labels had a great effect on how people interpreted a drawing. The clothing of a drawing, for instance, was interpreted as farm clothing when the label was "drawings by farm children", but as a city

uniform when the label was "drawings by city children." (For closely related findings see Chapman & Chapman, 1967; Medin & Wattenmaker, 1987).

Hampton (1987) provides another powerful demonstration from the domain of conceptual combination. He noted that some novel conceptual combinations give rise to *emergent properties* - properties that are not true of either constituent. For example, when asked to describe the novel concept "beach bicycle", participants listed attributes such as *has wide wheels*. Since "beach bicycle" was a novel category, such emergent properties cannot have arisen from consideration of category examples. Rather such emergent attributes must have resulted from people's relevant background knowledge. People, for instance, may have considered that riding on the beach with a normal bike would be problematic since its tires would sink in the sand. Naïve knowledge of physics dictates that such a problem can be remedied with flat tires (for a review on conceptual combination see Hampton, 1997).

It seems therefore that concepts are intrinsically intertwined with our 'intuitive theories', with how we make sense of the world. This view is commonly known as the *concepts-in-theories view* (also known as the theory-based or explanation-based view) and is eloquently voiced in Murphy and Medin (1985). 'Intuitive theories' refer to any host of mental explanations and may include anything from scientific principles to fallacious beliefs.

In terms of representational structure, the theory-based view can be safely interpreted to support the idea that concepts involve features embedded in networks of knowledge. A corollary of the theory-based view is that features, depending on the role they play in how we understand a concept, differ in how much coherence they lend to that concept. The more central a role a feature plays in our intuitive understanding of a concept, the more conceptually central that feature is. For the concept elephant, for example, the feature *heart* promotes more coherence than the feature *gray*. Thus *heart* is a more conceptually central feature for elephant than the feature *gray*.

To sum up, a place to search for determinants of category-based inference is by looking at (proposed) factors that constrain representational structure. Both conceptual and empirical evidence can be said to corroborate that 'intuitive theories' constrain conceptual

structure and use. Hence, it is reasonable to expect 'intuitive theories' to constrain property induction also. The present research can be seen as an extension of this corollary of the theory-based view, concerning the view that features differ in terms of conceptual centrality. In this thesis, the representation and focus of the theory-based view is extended to category-based inference. If conceptual centrality is a structural aspect of representations, then it should also influence property induction. Specifically, the more central a feature is (or is believed to be) to a target concept, the more projectible it should be to that concept. So our broad answer to the question "Why are some properties more projectible than others?" is "Because some properties (are believed to) promote more coherence to the target concept than others".

1.6 CONCEPTUAL CENTRALITY AS MUTABILITY

The claim that features differ in conceptual centrality depending on how central a role they play in how we understand a concept is appealing. At the same time, the notion of an 'intuitive theory', and thus the notion of a feature's centrality in such a theory, is underspecified. What is needed therefore is a clear articulation of conceptual centrality, one that will allow us to operationalize it and thus make it amenable to empirical testing. Such an articulation was offered by Sloman et al. (1998).

For a given concept, feature centrality can be defined in terms of how mentally transformable a feature is while retaining the integrity (coherence) of that concept (Sloman et al., 1998). Sloman et al. name this measure of conceptual centrality *mutability*, a term they borrow from Kahneman and Miller (1986). A feature is immutable (conceptually central) to the extent that presence of that feature in an object increases people's certainty that the object is represented by the concept.

A way to measure mutability is by asking people how easy they find it to transform a concept into one that is in all ways like the original but lacks that feature. This is known as the *Ease-of-Imagining Task* (see e.g. Sloman et al., 1998). For the concept dog, for instance,

the feature *heart* is more immutable (more conceptually central) than the feature *fur*, in that it is harder to imagine a dog without a *heart* rather than one without *fur*.

Our hypothesis is that mutability should constrain projectibility. People should be more willing, for instance, to project *heart* rather than *fur* to a newly discovered species of dog. Before we proceed though, it cries to be shown that mutability is not a superfluous concept - it is not reducible to known dimensions of category structure. In the first part of the following section, we provide conceptual arguments distinguishing mutability from variability, diagnosticity, and salience. In the second part of that section, we provide empirical evidence that validates these distinctions.

1.7 MUTABILITY VERSUS VARIABILITY, DIAGNOSTICITY, AND SALIENCE

1.7.1 Conceptual arguments

Mutability is not variability. Following Tversky and Kahneman (1983), we assume that there are two distinct perspectives that people may adopt about a concept, an *inside* and an *outside* view. The inside view looks at the mental representations of categories such as features, and elements that bind these features together. The outside view looks at sets of actual or remembered instances that comprise a category. Mutability concerns an inside view, since it measures the extent to which a feature is mentally transformable while retaining the integrity of a concept. In contrast, variability involves an outside view since it measures the extent to which a feature is transformable across the remembered instances of a category.

Most of the times the inside and outside views of concepts are compatible, and hence measures of mutability and variability converge. For example, *heart* is both an immutable and a homogeneous property of mammal. In other cases though, mutability and variability measures may diverge; e.g., the color of ravens is one of their very homogeneous properties but still one of their mutable properties. It seems relatively effortless to imagine a raven that is not black.

Mutability is not diagnosticity. There are two predominant uses of the term diagnosticity, one in terms of *informational value* and one in terms of *inferential potency*. In terms of informational value, a feature is diagnostic to the extent that it helps differentiate an object belonging to one category versus others. Being striped, for instance, is a highly diagnostic feature of zebras relative to the set of animals in the sense that a striped animal is highly likely to be a zebra. Conceptually, mutability and informational value are easily distinguishable. *Having a liver*, for instance, is an immutable property for zebra though it has a very low informational value since many things have livers but are not zebras.

In the latter sense of the term, diagnosticity is equated to the inferential potency of a feature. Inferential potency refers to the extent to which knowing that an object has a feature helps predict that the object is also likely to have certain other features as well (e.g., Franks, 1995). Mutability is conceptually distinct from this second sense of diagnosticity. A feature derives inferential potency by virtue of its correlation to other features. In contrast, a feature derives its mutability status by virtue of dependencies. For instance, *has buttons*, is statistically correlated with *has a zipper*, *has material*, and *is colored* (Malt & Smith, 1984) and thus an inferentially potent feature. That is, knowing that an object has buttons allows one with relative confidence to predict that the object is likely to also have a zipper, material, and color. At the same time *has buttons* is a relatively mutable feature; it is easy to imagine, for instance, a shirt that does not have buttons.

Mutability is not salience. Salience measures the extent to which a property pops-out from the background - the extent to which it is noticeable in a context-independent sense. A loud noise, for instance, is more salient than a barely audible one. Clearly salience is distinct from mutability; the florescent jackets of policemen is one of their salient properties albeit a mutable one.

To sum up, mutability can be conceptually distinguished from other dimensions such as variability, diagnosticity, and salience.

1.7.2 Empirical evidence

Evidence supporting that mutability is distinct from variability, diagnosticity and salience is presented in Sloman et al. (1998). In Study 1, they asked participants to provide estimates for various measures. Some of these measures were believed to tap mutability, some variability, some diagnosticity, and some salience. All the measures of mutability used had in common that they asked participants to consider an object lacking a feature but otherwise intact. A factor analysis of the results provided evidence that (i) all the measures believed to tap mutability indeed gauged the same underlying factor, and (ii) that this factor was distinct from the underlying factors of diagnosticity and salience measures. That particular study failed to differentiate mutability from variability. Study 5 pitted mutability against variability and supported that these two dimensions are empirically distinguishable (for further evidence see also Ahn & Sloman, 1997).

To sum up, both conceptual arguments and empirical evidence support that mutability is psychologically real in two senses: (i) in that different measures of mutability gauge on the same underlying factor, and (ii) that this factor is distinguishable from the factors of variability, diagnosticity, and salience. Having thus established that it is constructive to use mutability as an explanatory construct, Sloman et al. (1998) went on to model mutability as centrality in a network of pair-wise dependency relations.

1.8 *MUTABILITY AS CENTRALITY IN A NETWORK OF PAIR-WISE DEPENDENCY RELATIONS*

Under the simplifying assumption that representations involve features embedded in networks of relations which are generic (their type is irrelevant) and asymmetric (non-reflexive), Sloman et al. (1998) modeled mutability as centrality in a network of pair-wise dependency relations. According to their model, a feature is central, for a particular concept, to the extent that other (central) features depend on it. The main idea behind this claim is that

changing features upon which lots of other features depend, should destroy the integrity of a concept because it would force a cascade of other changes. In contrast, changing features that are peripheral in a concept's dependency structure, should leave the structure of a concept relatively unaffected. Applying this notion to the example about elephants, given that a person believes that lots of features about elephants depend on *heart* but only few on *gray*, the former feature is more conceptually central than the latter. In most of the present experiments, we manipulate the relative centrality of features by manipulating their centrality in a concept's dependency space.

Sloman et al. implemented their centrality hypothesis with the simple, iterative, linear equation:

$$c_{i,t+1} = \sum_j d_{ij} c_{j,t}$$

where $c_{i,t+1}$ is the centrality of feature i at time $t+1$
 d_{ij} is the extent to which feature j depends on feature i ,
and $c_{j,t}$ is the centrality of feature j at time t .

To test their model, they presented participants with features of particular categories and, for each category, they asked them to draw links binding these features together (Study 2). Participants were given the choice of three colors to indicate the strength of a dependency. These dependency estimates provided information about the d_{ij} terms of the model, which is the only information the model needs to predict the relative centrality of properties within a category - the model has no free parameters. The results from the model were compared to mutability estimates for the same features. The model provided reasonably good fits of the mutability data. Correlation measures between the model's predictions and mutability estimates were in the predicted direction for all categories. Also, the mean correlation between the model's predictions and the mutability data was better than .61 for all categories. To summarize, the centrality model provides reasonably good fits to the data given that the

model has no free parameters. This supports the notion that conceptual centrality defined as centrality in a concept's dependency space is a good surrogate of mutability.

The Sloman et al. (1998) model can be said to act as a bridge between the appealing but not very well articulated concepts-in-theories approach, and the less appealing but better articulated feature-based approach. The asymmetric dependency links connecting features in the centrality model can be interpreted as stand-ins for intuitive micro-theories. The knowledge that wings enable birds to fly, for example, is reductively represented in the model as *can fly* → *having wings*, where → denotes a generic, asymmetric, dependency link. Following this interpretation, the model's conceptual centrality estimate for a given feature can be viewed as a global measure of the importance of a feature in a person's 'intuitive theory' about a concept. Notice that the Sloman et al. model is not necessarily confined to 'intuitive theories' - it can represent non-explanatory connections. For instance, with its unlabeled asymmetric dependency links it can perhaps also capture pre-theoretical biases such as those reported by Keil (1995).

1.9 CONCEPTUAL CENTRALITY AND CATEGORY-BASED INFERENCE

Although Sloman et al. provide a model that quite successfully captures the relative centrality of features within a concept, the present focus is on the role of centrality in the projectibility of properties across concepts. This section leads to our proposal of how feature centrality influences property induction.

How does centrality constrain projectibility? As a starter, one might propose that the more conceptually central a feature the higher its projectibility. Name this Hypothesis 0.

Hypothesis 0: all else being equal, the more conceptually central a feature, the higher its projectibility across concepts.

Hypothesis 0 captures the idea that conceptually central features should be weighted more heavily in inference. That is, central features should be more projectible than less-central ones, because projecting a central feature should enhance more the coherence of the conclusion (or target) category. What Hypothesis 0 fails to capture is that centrality is *concept relative*, i.e. one and the same feature can have different mutability statuses across different concepts. For example, *roundness* is a central feature of basketballs but not of cantaloupes (cf. Medin & Shoben, 1988). In sum, Hypothesis 0 seems (at best) to be confined to cases of generalizations from instances (e.g., an inference from a sample of bananas to all bananas) but it fails to capture category-based inference (e.g., an inference from bananas to boomerangs).

The observation that centrality is concept relative is readily accounted for by the Sloman et al. definition of centrality. Different concepts, such as bananas and boomerangs, generally involve different dependency structures and hence the centrality status of a given feature across concepts may well be different. For example, lots of things about boomerangs depend on the fact that they are curved (e.g., their ability to return when thrown), whereas almost nothing about bananas depends on that property. Bananas are something we eat and hence *being curved* is peripheral to their dependency structure.

The present research focuses on category-based inference, and hence Hypothesis 0 needs to be revised. Here is the hypothesis we propose:

The Centrality Hypothesis: All else being equal, the more central a feature for a concept, the higher its projectibility to other concepts to the extent that those concepts share its structure.

Our centrality hypothesis derives from an appeal to conceptual centrality and a belief that projecting a property aims to maximize coherence. The more the premise and conclusion categories share structure, the more likely it is that a property will retain a similar centrality status in the target concept. The more central a property is believed to be in the (dependency) structure of the target concept, the more coherence its projection would promote.

Our appeal to conceptual centrality and our view that projecting a property aims to maximize coherence, predicts property by category effects. For instance, even though *color* is a relatively non-projectible attribute, it should lead to strong inferences from solar panels to solar calculators. Naïve knowledge of physics informs us that lots of the functions of solar panels (such as attracting and storing energy) depend on their color being dark. Because solar calculators share a similar dependency structure as solar panels with respect to these functions, the argument should be judged strong.

Our line of reasoning is quite similar to Quine's (1961, 1977) argument against the existence of *analytic truths*, statements necessarily true by virtue of the language. A whole approach to concepts, commonly known as *the classical view*, is based on the assumption that concepts have defining properties. It holds that some properties, such as *we sit on it* for the concept chair, are defining, and thus statements like "chairs are objects that we sit on" are analytic truths. Quine argued that although almost any feature can be removed from a category, we are less willing to give up features upon which much of our knowledge depends. That is because deletion of such features will propagate lots of other revisions in our knowledge base. Quine argued therefore for a continuum, rather than a dichotomy, from analytic to synthetic truths.

In a similar vein, we hold that features can be arranged in a continuum of conceptual centrality reflecting how much of the structure of a concept depends on them. Further, we hold that central features should be more generalizable than non-central ones across concepts that share similar structure, because projecting central features across such concepts should promote more coherence in the target concept.

1.9.1 What we do not claim

To clarify our hypothesis, we stress what we do not claim. First of all, we make no claims about how 'intuitive theories' (or pre-theoretical biases) come about - we take them as givens (but see, for example, Atran, 1995, 1998; Carey, 1995; Keil, 1995; Keil, Smith, Simons, & Levin, 1998). Further, we do not assume that representations are stable across

individuals. Quite the contrary, we believe that representations (and thus feature centrality) are conditional on a person's knowledge. People have different theories about how the world works, and hence their category representations should involve different dependency structures. To illustrate, for most non-biologists like myself, the fact that ravens are black is a relatively non-central property - we have no theories connecting raven's color to the rest of their properties and functions. Ornithologists, on the other hand, possibly have elaborate theories of why ravens are black, theories that relate ravens' blackness to the rest of their features. For instance, they may know that their color regulates mating, increases body temperature, facilitates enzyme activity and so on (these explanations are made up; their sole aim is to clarify the argument). What we believe remains stable across individuals is the way conceptual centrality influences inference. In other words, we believe that centrality provides a structural constraint on category-based inference. Finally, we have no illusions of trying to offer an all-encompassing theory of what constrains categorical inference. We merely suggest that the extent to which a feature is believed to lend coherence to a concept should affect its projectibility to other similar concepts.

--OBJECTIVES AND DESIGN OF THE PRESENT RESEARCH--

1.10 OBJECTIVE 1: ASSESSING THE CENTRALITY HYPOTHESIS

The main objective of the current research is to investigate empirically the hypothesis that central features are more projectible than less-central ones across categories that share similar structure. There are two parts to this claim: (i) that conceptual centrality influences the inductive potential of features, and (ii) that the effect of centrality on projectibility is mediated by the extent to which the premise and conclusion concepts share similar dependency structure. To empirically assess our hypothesis we manipulate (i) the (relative) conceptual centrality of features, and (ii) the similarity between an argument's premise and conclusion categories.

1.10.1 Manipulations of conceptual centrality

Our basic experimental paradigm involves informing participants that a premise (or base) category has two novel properties, a central and a less-central one. Participants are then asked to rate the likelihood that a conclusion (or a target) category has each of these properties, or in a few cases to choose which of the two properties the conclusion category was more likely to have. The relative centrality of the properties for a given argument was either assessed by tasks known to tap mutability (for a detailed exposition of such tasks see Sloman et al., 1998), or else it was directly manipulated based on the Sloman et al. rough definition of centrality according to which a feature is central to the extent that other (central) features depend on it.

The *Ease-of-imagining* task. The ease-of-imagining task measures the mutability of a feature by asking people how easily they find it to imagine an object missing that feature which is otherwise intact. To assess, for instance, the relative centrality of the features *heart* and *gray* for the concept elephant, participants are asked to imagine an elephant and then to transform this representation into one that lacks the relevant property. The more difficult participants find it to complete the transformation, the more conceptually central the feature is taken to be. In the example above, participants should find it more difficult to imagine an elephant lacking a *heart*. Mutability estimates are then compared to inductive strength estimates about the properties in question to assess whether our centrality hypothesis is supported. Experiments 1 and 2 of Chapter II assessed conceptual centrality by the ease-of-imagining task.

Centrality from a dependency chain. Experiments 3 to 7 of Chapter III, define the relative centrality of two properties by embedding them in a dependency chain: the central property being the one upon which the less-central property depends (either directly or via a mediating property). This operationalization of centrality is based on the Sloman et al. notion that centrality can ultimately be computed from local dependency relations. An example of an argument using such a manipulation of centrality is,

Lions have both the enzyme amylase and the enzyme transacetylase.
For lions, the enzyme transacetylase is regulated by metabolism.
For lions, the enzyme amylase regulates metabolism.

- A. Tigers have transacetylase. _____ %
B. Tigers have amylase. _____ %

The enzyme amylase is relatively central since metabolism depends on it. The enzyme transacetylase is not as central, since it depends on metabolism.

This operational definition of centrality aimed in part to pit our theory's predictions against those of structural-mapping models of inductive inference. This issue will be fully exposed in Chapter III.

Centrality from the number of the dependent properties. A more natural way to operationalize relative feature centrality while staying close to the spirit of the Sloman et al. (1998) definition, is by manipulating the number of properties that depend upon a feature. This is how we manipulate centrality with the "Lots/Few" definition in Experiments 8 to 11 of Chapter IV (animal categories), and Experiment 17 of Chapter VII (artifacts). We stipulate that members of a category have two properties: one upon which lots of their functions depend, and one upon which only few of their functions depend. Central properties are assumed to be the ones upon which lots of the category's functions depend. Below is an example of an argument using the "Lots/Few" definition of conceptual centrality,

Dolphins have the neurotransmitter oxytocin and the neurotransmitter calcitocin.
Lots of a dolphin's functions depend on the neurotransmitter oxytocin.
Few of a dolphin's functions depend on the neurotransmitter calcitocin.

- A. Seals have calcitocin. _____ %
B. Seals have oxytocin. _____ %

In this argument, the central property is the *neurotransmitter oxytocin*, and the less central property the *neurotransmitter calcitocin*. Because dolphins and seals are very similar, we expect participants to give a higher rating for conclusion B.

Experiments 15 and 16 in Chapter VI asked what sort of information (centrality or else variability) people seek to make a categorical inference. Participants had to guess whether a target species had a property, based on the information that some members of another species (the base) had that property. To facilitate their decision, they could either seek conceptual centrality or else variability information. The conceptual centrality information, as in the "Few/Lots" definition, concerned how many of the base category's features and functions depended on the critical property. The variability information concerned the proportion of the members of the base category that had the critical property. To the extent that centrality influences more category-based inference than variability, we expected participants to choose centrality information.

Centrality from the centrality of the dependent properties. A subtle way to manipulate the relative centrality of two features is to keep the number of properties that depend upon them constant, and manipulate the centrality status of the properties that depend upon them. Recall, that a feature is central to the extent that other central features depend on it. Experiments 12 to 14 reported in Chapter V use this definition of centrality. As an example, we stipulated that hippos have the hormone aldosterone and the hormone corticosterone. We also stipulated that, for hippos, aldosterone regulates the amount of oxygen in the blood, controls blood density and blood pressure, whereas corticosterone regulates the amount of green in the eye, eye-colour reflectance and uniformity of eye-colour. Because the properties that depend upon aldosterone are (intuitively) more central than those depending upon corticosterone, we assumed that aldosterone was more central than corticosterone.

In sum, various measures of conceptual centrality were used to manipulate relative feature centrality. There are good reasons supporting the use of many centrality measures. For one, it helped us to pit our hypothesis against others that were supported by using comparable manipulations. Secondly, as will become evident in the experimental chapters, the results obtained by each definition on its own were open to various interpretations. Using a variety of centrality measures helped to eliminate competing interpretations across sets of experiments. That is, the use of multiple operational definitions strengthens the force of our centrality

hypothesis. Finally, a future aim is to model centrality effects on inductive inference. The Sloman et al. (1998) notion of conceptual centrality comes with a model that implements it. Thus, it provides a good basis for building category-based models to account for centrality effects. Using multiple ways of operationalizing centrality from the Sloman et al. definition, provides therefore valuable evidence for constructing such models. In the Final Discussion we will discuss such models.

1.10.2 Manipulations of shared structure

The second part of our hypothesis states that the effects of centrality on category-based inference should depend on the extent to which the premise and conclusion categories share structure (features and dependencies). This is based on the assumption that the more structure two concepts share, the more likely the candidate property is to share a similar centrality status in the target concept. The more central a status a property is believed to have in the target, the more coherence projection of the property would promote. Even though we expected the effect of centrality to be bounded by the amount of premise-conclusion overlap, we had no *a priori* ideas about the exact amount (or type) of shared structure required. This issue was addressed empirically by manipulating the extent to which premise-conclusion category pairs were judged similar. Our working assumption was that the more similar a pair of categories is judged, the more dependency structure they share. The straightforward prediction therefore was that the more similar a pair of categories was judged, the bigger the difference between the projectibility of the central versus the less central feature. Sometimes we selected category pairs that have been reported in the literature to be associated with high or low similarity estimates. In most cases we also independently checked the assignment of category pairs to similarity conditions by gathering similarity estimates from a separate group of participants.

Some of our 'natural kind' experiments manipulated shared structure by using high or low-similar mammal pairs (e.g. hippos-elephants, mice-elephants). Other experiments manipulated shared structure more widely by using similar animal pairs from the same

superordinate (e.g. tigers-lions), dissimilar pairs from the same superordinate (e.g. chimps-lions), and dissimilar pairs from a different superordinate (e.g. falcons-lions). Pairs of the first shared structure condition were judged as more similar than pairs of the second condition which were judged more similar than pairs of the third condition. We did not include pairs from different superordinates but high in similarity simply because it is very hard to find any (bats being similar to some bird species, and seals and dolphins being similar to fish are some of a very limited set of such examples). As a result, our experiments do not fully address the impact of superordinate structure on inductive reasoning (recall that Osherson et al. (1990) assume that categories are hierarchically structured and thus reasoning recurses to taxonomies, though Sloman (1997) provided evidence against this claim). Our experiments, however, do address two more impoverished claims: (1) that sharing a superordinate decides the strength of arguments, and (2) that sharing a superordinate mediates the centrality effect on property induction. If claim (1) were true, then we should not observe differences in property generalizability among base-target categories from the same superordinate, but we should observe differences for categories from a different superordinate. Similarly, if (2) were true, then we should not observe differences in the generalizability of the central and the less central properties among same superordinate categories, but we should observe differences among different superordinate categories.

Dependent properties: abstract versus specific. To better understand what type of shared structure matters in projecting central versus less central features, some experiments left the properties that depended upon the candidate feature abstract (like the experiments using the "Few/Lots" definition), while others specified them (like most experiments that defined centrality from a dependency chain). Specifically, we investigated the possibility that feature projectibility depends on the extent to which the target category shares (or is believed to share) the dependent properties specified by the candidate property for the base. Elephants and mice, for instance, are fairly dissimilar mammals though they are similar in terms of specific properties such as fur, and heart.

In addition, the use of abstract dependent properties helped assess whether feature centrality operates under vagueness. To the same end, we also used abstract predicates, such as *has substance Z*, and poorly defined categories, such as Mammal 1 and Bird 2. Note that for some models of category-based inference to operate, such as structural-mapping models, they require a complete specification of the candidate properties as well as the properties that depend upon them. Finding that the feature centrality hypothesis operates under conditions of vagueness, cannot be accounted for by such models. Further, such a finding would give the centrality hypothesis a heuristic status.

In sum, the overarching objective of our research is to examine empirically the hypothesis that central features are more projectible than less-central features across concepts that share similar structure. To assess our hypothesis empirically we needed to operationally define 'feature centrality' and 'shared structure'. To this end, we stated how centrality and shared structure was measured or manipulated. Further, we provided reasons for using the various operational definitions of feature centrality and shared structure.

There are two other related objectives that our research pursues, and are reviewed in the following sections.

1.11 OBJECTIVE 2: MUTABILITY VS. VARIABILITY

One of our working assumptions is that mutability and variability involve different perspectives for looking at categories. We take mutability to be an aspect of mental representations in an 'inside view' of categories, whereas variability to be an aspect of remembered instances in an 'outside view' of categories. A reason to believe that this distinction is grounded comes from the work of Medin and Shoben (1988). They provided empirical evidence that properties judged as equally typical for two categories, can nonetheless have a different mutability status across these categories. Although, for instance, roundness is an equally typical property of the categories wheel and orange (almost all the

wheels and oranges one can remember are round), deleting the property *round* plays more havoc in the representation of wheels than that of oranges. Our straightforward prediction was that the effect of conceptual centrality and variability on the projectibility of a property should be distinct. For example, we expected people to be more confident to endorse the property *round* to an unseen wheel rather than to an unseen orange. That is because projecting the property *round* to an unseen wheel should promote more coherence than projecting this property to an unseen orange.

Most of our experiments controlled for the variability of the central and the less-central properties; the instances of the base category were assumed to have both properties. Moreover, the experiments of Chapters V and VI directly pitted centrality against variability to examine the relative effect of each on projectibility. Experiments in Chapter V manipulated the frequency of the central and the less-central properties, such as the less-central properties were more homogeneous than the central ones. We expected feature centrality to provide a stronger constraint on projectibility. The experiments in Chapter VI, pitted mutability against variability in terms of the information people seek to make an inference decision. Participants' task was to guess whether a target species had a property, based on the information that some members of another species (the base) had that property. To facilitate their decision, they could either seek conceptual centrality information (information about how many of the base category's features and functions depended on the critical property) or else variability information (information about the percentage of the members of a category that had the critical property). To the extent that centrality influences more category-based inference than variability, we expected participants to choose centrality information.

1.12 OBJECTIVE 3: INFERENTIAL PROCESS: DOMAIN GENERAL OR DOMAIN SPECIFIC?

A final goal of the present research is to show centrality effects on category-based inference in different conceptual domains. Previous work on property induction mostly

focused on the natural kind domain (though exceptions are provided in Shafir, Smith, & Osherson, 1990, and in Sloman, 1993, 1998). At the same time, most models of psychological strength (such as the similarity-coverage model and the feature-coverage model) suppose domain-general processes. It is transparent that this claim needs to be assessed by using arguments beyond the natural kind domain.

The need to consider arguments from different domains is even more pressing because some theorists entertain the possibility that inference procedures depend on domain-specific theories (like domain-specific causal laws, Carey, 1985). Gelman (1988) claims that natural kinds have more inductive potential than artifacts. In a similar vein, Keil (1989) points to many differences between the artifact and natural kind categories. His general claim is that properties of artifacts are projectible to the extent they are related to their purpose or design.

Our research assumes that effects of centrality should be domain-general. Artifact categories should also have rich dependency structures, and hence centrality should be an aspect of artifact representations as well. We know, for instance, that the color of solar panels helps them to attract and store solar energy. Both of these functions are crucial in how we conceive of solar panels. Hence, the color of solar panels seems also to be one of their relative central properties. To be clear, we do not claim that there are no general differences between different domains. Nonetheless, we suppose that projectibility should not depend on type of category but rather on the relation between predicate and category (Goodman, 1955, Markman, 1989, Sloman, 1998).

The present research addressed the domain-general versus the domain specific issue by using categories from both the natural kind and artifact domains. Specifically, the experiments reported in Chapters II to VI involved animal categories, whereas the one reported in Chapter VII involved artifacts.

1.13 OVERVIEW OF THE EXPERIMENTAL CHAPTERS

The experiments are organized into six experimental chapters, roughly reflecting the operational definition of conceptual centrality used. In Chapter II mutability was measured from the ease-of-imagining task, in Chapter III from a single dependency chain, in Chapters IV and VII by the number of functions that depended upon a property, and in most studies of Chapter V by the centrality of the properties that depended upon the candidate property (refer to the section 1.10.1 Manipulations of conceptual centrality). The sets of studies in Chapters II and V aimed to show that conceptual centrality influences projectibility. They also aimed to detect a boundary condition for centrality effects by manipulating (at various levels) the amount of shared structure between the base and target categories. In other words, these experiments addressed our conceptual centrality hypothesis.

Experiments of Chapter V pitted centrality against variability (the frequency of a property across instances of a category) to detect which dimension exerts more influence on projectibility. Experiments of Chapter VI took a step back and asked what sorts of information (centrality or else variability) people seek to make a categorical inference. Finally, the experiment of Chapter VII aimed to extend the centrality hypothesis to the domain of artifacts.

Below we summarize the broad issues investigated by the present experiments, along with the chapter(s) that addressed them:

1. Role of centrality in projectibility (all chapters)
2. Role of shared structure on centrality effects (all chapters except Chapter II)
3. Domain generality of the centrality hypothesis (addressed by relating the results of Chapters II to VI that used animal categories with those of Chapter VII that used artifact categories).
4. Role of vagueness (mostly Chapters IV, VI, and VII)
5. Comparative roles of mutability and variability on projectibility (addressed in Chapters V and VI).

CHAPTER II:

TWO MEASURES OF CONCEPTUAL CENTRALITY: IMMUTABILITY AND CENTRALITY IN A DEPENDENCY SPACE

The main objective of our project is to examine the hypothesis that the extent to which a feature is central for a concept, the extent to which a feature lends coherence to a concept, affects its generalizability to other concepts. To do so we need tasks that measure conceptual centrality. A subclass of such tasks involves asking participants to perform conceptual transformations. Participants are asked, for instance, to mentally transform a concept into one that is in all ways like the original one but lacks the critical feature (for a variety of conceptual centrality measures see Sloman et al., 1998). The harder people find it to perform the mental transformation, the more immutable the feature is assumed to be. For example, *has wings* is more immutable than *has a tail* for robin, in that people should find it easier to imagine a robin without a tail than one without wings.

Experiment 1 sought empirical evidence for the centrality hypothesis by measuring feature centrality from a conceptual transformation task. Participants were presented with pairs of category-based arguments. Both arguments in a pair had the same premise and conclusion categories but involved different candidate properties. Participants were asked to choose the argument they deemed more convincing. Participants were also asked to rate the immutability of a property from the ease-of-imagining task. Our straightforward prediction was that judgments of immutability should be an increasing function of argument strength estimates.

Showing that immutability influences projectibility provides but a first step in our research. Ultimately we aim for a clearer articulation of conceptual centrality, an articulation that will allow us to develop precise empirically testable hypotheses (and ultimately a model) of how feature centrality affects category-based inference. Such an articulation was provided by Sloman et al. (1998). Under the assumption that concepts involve features embedded in networks of asymmetric dependency relations, Sloman et al. claimed (and provided empirical support for this claim) that immutability can be surrogated by centrality in a concept's dependency space. Roughly, the more properties and functions of a concept depend upon a candidate feature, the more conceptually central that feature is. Applying this notion to the

example about robins, *has wings* is more conceptually central than *has a tail* because more properties and functions about robins depend on having wings than on having a tail.

Experiment 2 provided a first step toward applying the Sloman et al. notion of centrality to category-based inference. The design was quite similar to that of Experiment 1 with the change that participants were asked to provide conceptual centrality estimates from both the ease-of-imagining task and the dependency task. The dependency task asked participants to estimate how much about the structure of a concept depended on the candidate property. The higher the estimate, the more conceptually central a property was assumed to be. The straightforward prediction was that conceptual centrality measured from the dependency task should also be correlated with argument strength. Further, to the extent that the dependency task and the ease-of-imagining task hone down to the same factor (i.e. conceptual centrality), we expected no response differences based on the task used.

In sum, the present experiments attempted to show that feature centrality, as either immutability or else centrality in a dependency space, constrains property induction. Before we proceed with the experiments, we motivate our centrality hypothesis by considering evidence from the closely related domain of categorization as well as evidence from the domain of category-based inference. We also take a closer look at the design of the experiments.

2.1 MOTIVATION

In the introductory chapter we motivated our centrality hypothesis by providing evidence that lay theories constrain what features get tabulated and represented. We reasoned that constraints on conceptual representations should also constrain conceptual use. Below we motivate our centrality hypothesis more directly by considering evidence from the domains of categorization and category-based inference.

2.1.1 Categorization: Differential weighting of deep vs. surface features

One reason to expect that feature immutability influences projectibility comes from studies on categorization. Keil (1989), for instance, found that children prefer to categorize a raccoon that was painted black with a white stripe on its back as a raccoon rather than as a skunk. Similarly, Rips (1989) found that adults prefer to categorize a bird that was chemically contaminated and came to look like an insect but could still mate with birds as a bird rather than an insect. These studies show that deep features, such as *mates with birds*, are weightier in category decisions than surface ones.

For animal categories, deep features can be said to be more conceptually central than surface ones by virtue of their dependencies. Surface features depend on deep features while deep features do not necessarily depend on surface ones. An animal's phenotype, for instance, depends upon an animal's genotype, while an animal's genotype does not depend on the animal's phenotype. The color of a raccoon, for instance, depends on its genetic structure, while the raccoon's genetic structure does not depend on its color.

Categorization and inductive inference are intimately related: inductive inference is a major function of categorization (e.g. Smith, 1990). Having categorized an entity as a bird, for example, helps infer with (relatively) high confidence that it will have wings, it will lay eggs, it will be feathered and so on. Moreover, knowing that instances of a category have a property, sanctions the inference that members of similar categories will have the property as well. Knowing, for example, that robins have the neurotransmitter dihedron makes it likely that sparrows (but possibly not spiders) will have this neurotransmitter as well. To the degree therefore that conceptual centrality influences categorization, we also expect it to influence category-based inference.

2.1.2 Category-based inference: Projectible vs. nonprojectible properties

Our research is also motivated directly from studies on category-based inference (e.g. Carey, 1985; Gelman, 1988). These studies assumed that some properties were more generalizable than others, and tested for property effects on category-based inference. We

argue that some of these property effects can be explained by appealing to conceptual centrality.

Gelman (1988) presented preschool and elementary school children with various facts about a (base) category (e.g. "this rabbit has a spleen inside") and asked them to report whether these facts generalized to other (target) categories (e.g. a brown rabbit, a dog, or a phone). One of the factors that Gelman manipulated was the type of the candidate property (generalizable or else non-generalizable). Generalizable properties involved properties that seem stable within an individual (such as *has pectin inside*), while non-generalizable properties ones that seem variable (such as *is cold*). Gelman found that both preschool and elementary school children projected generalizable properties in proportion to the similarity between the target and the base categories. For non-generalizable properties children's responses were at chance level - they were unaffected by similarity condition.

Some properties may seem relatively stable because they are conceptually central. Other properties may seem relatively unstable because they are not conceptually central. More of a rabbit's properties, for instance, depend on properties such as *has pectin inside*, than on properties such as *is cold*. To the extent therefore that properties that seem intrinsic or stable are more conceptually central than properties that seem extrinsic or unstable, Gelman's results corroborate our centrality hypothesis.

2.2 BASIC EXPERIMENTAL DESIGN

Chapter II reports two experiments that aim to show that conceptual centrality constrains projectibility. Experiment 1 asked participants to choose between pairs of categorical arguments that shared premise and conclusion categories but involved different properties. Participants were asked to choose the argument that they found more convincing, the one whose premise better supported its conclusion. This estimate should tap directly on argument (or psychological) strength. Participants were also asked to provide conceptual centrality estimates for each property from the ease-of-imagining task. The rationale was that

by analyzing conceptual centrality and argument strength estimates we could detect whether centrality influences projectibility. Our prediction was that participants would rate as stronger the argument whose premise feature they rate as more conceptually central.

Experiment 2 used the same materials as Experiment 1. Pairs of arguments used in Experiment 1 were split apart into two separate arguments. Participants had to provide conditional likelihood estimates for each resulting argument. Participants in Experiment 2 therefore made an independent estimate of the likelihood of each individual argument. Experiment 2 aimed to rule out the claim that the results of Experiment 1 are due to a demand characteristic stemming from its forced choice design.

In addition, each participant was asked to provide unconditional likelihood estimates for the premise and conclusion statements of each argument. Note that conditional likelihood estimates may be contaminated by prior beliefs in an argument's conclusion. By regressing conceptual centrality estimates onto conditional likelihood estimates while partialling out the effects of prior knowledge we aimed to detect whether mutability influences argument strength.

Experiment 2 used two tasks to measure conceptual centrality: (i) the ease-of-imagining task, and (ii) what we call the dependency task. The dependency task asked participants to assume that a category had the candidate property and to estimate how many properties about that category depended upon this property. This task taps directly on the Sloman et al. (1998) notion that immutability can be surrogated by centrality in a concept's dependency space. The more properties and functions of an animal participants believed depended upon the candidate property, the more immutable the property was considered to be. To the extent that both tasks measure conceptual centrality, we expected the effect of centrality on projectibility to be independent of the task used.

2.3 EXPERIMENT 1: IMMUTABILITY

METHOD

Participants. Twenty-four first-year psychology undergraduates of the University of Durham volunteered to participate.

Design and Materials. Participants were given two tasks to complete: the Induction task and the Centrality Estimation task.

Induction task. Participants were given 12 items to evaluate, each one consisting of a pair of categorical arguments. Both arguments of an item had the same premise and conclusion categories, but involved different properties. Participants were asked to circle the categorical argument (A. or else B. in the example below) that they deemed more convincing. A sample item from the induction task was:

-
- A. Fact: Giraffes have dense nerve fibres in the mid-brain.
Conclusion: Zebras have dense nerve fibres in the mid-brain.
- B. Fact: Giraffes have incisor teeth encircled by enamel.
Conclusion: Zebras have incisor teeth encircled by enamel.
-

The full set of materials used is presented in Table 2.1 (see overleaf). Note that the predicates used were relatively blank. We avoided using familiar predicates because such predicates would presumably contaminate participants' inductive choices.

Centrality Estimation task. The conceptual centrality of a property was estimated by the ease-of-imagining task. Participants were asked to estimate how easily they could imagine an animal missing a property that it normally has. They were given to evaluate 48 items: the 12 premise and 12 conclusion categories paired with each of the 2 properties. For the induction item shown overleaf, for instance, participants had to provide four centrality estimates: two for zebras (one about enamel and one about dense nerve fibres) and two for

giraffes (one about enamel and one about dense nerve fibres). The first of these four items is presented below:

Imagine a zebra that has all the characteristics and properties of a zebra. Now change this image of the zebra so that it is in all ways like a zebra except that it does not have dense nerve fibres in the mid-brain that zebras have. Rate the ease of this transformation.

1	2	3	4	5	6	7	8	9
very easy			neutral			very difficult		

Table 2. 1 The full set of materials of Experiment 1. The left panels of the table present the premise-conclusion categories and the right panels the properties used.

Premise-conclusion categories	Predicates
Dogs-Squirrels	Use serotonin as a neurotransmitter
Lions-Elephants	Have poor night vision
Hippos-Tigers	Have PPT in their blood
Weasels-Bears	Can sense earthquakes before they occur
Goats-Bulls	Have synaptic vesicles containing zinc
Buffaloes-Rhinos	Release noxious chemicals when captured
Rabbits-Ferrets	Secrete uric acid crystals
Deer-Antelopes	Have four molar teeth
Giraffes-Zebras	Have adrenergic receptors
Camels-Wolves	Have high amounts of melanin in the skin
Seals-Flamingoes	Produce THS by their pituitary gland
Rats-Raccoons	Change their teeth twice in their lifetime
	Have glycine in their spinal cord
	Have sparse hair in their outer ear
	Have high density brain cells
	Have strong peripheral vision
	Have dense fibres in the mid-brain
	Have incisor teeth encircled by enamel
	Have low blood pH
	Have rolling ankles
	Have receptors to maintain blood pressure
	Have external ear that lack cartilage
	Have a choroid membrane in the eyes
	Have a new underfur growing in autumn

For each participant, the centrality estimates were used to divide the members of argument pairs into "central" and "less central". The proportion of times the central member was chosen in the forced choice induction task was the dependent variable.

Task order was counterbalanced across participants. The position of categorical arguments in the induction task was also counterbalanced across participants. Thus, 4 lists of materials were used. The order of the items in each task was randomized for each participant.

Procedure. Participants were presented with a booklet, containing one of the resulting 4 lists of material. There was no time limit but participants were encouraged to work quickly.

Induction task. Participants were presented with the 12 pair of arguments for evaluation. They were informed that they would be presented with a series of items, each one consisting of a pair of arguments. Individual arguments involved a fact followed by a conclusion. Participants were asked to choose the argument that they found more convincing, "...which argument's Fact provides a better support for its Conclusion." They had to work through an example of the same format as the test items. Participants were also told not to be concerned if some terms seemed unfamiliar to them. They just had to choose the argument that seemed more convincing to the best of their ability. Participants were instructed to ask the experimenter in the case of queries.

Centrality Estimation task. Participants were presented with the 48 ease-of-imagining questions for evaluation. They were informed that they would be asked to imagine an animal that lacked a property which it normally has. Their task was to estimate how difficult it was to complete this transformation on a 9-point scale. A rating of "9" meant that they found extremely difficult to imagine the animal lacking the property, whereas a rating of "1" meant that found it very easy to complete the transformation. The higher the rating the more difficult it was to imagine the animal missing the property. Participants were given to work through an example. They were instructed to ask the experimenter in the case of a query. Participants were asked for estimates for both the premise and conclusion categories.

RESULTS

Centrality Estimation task. For each participant, the member of each argument pair that had the highest centrality score was designated as "central", the other as "less central".

Thus, across all participants, the "central" and "less central" items were not the same. Table 2.2 presents the mean centrality scores for the central and the less central items for each task order. Unsurprisingly, the mean centrality scores for central items are higher than the corresponding scores for the less central items. The mean scores for the group that received the centrality task first are lower than the corresponding scores for the group that received the induction task first.

Table 2.2 Mean (SE) centrality estimates for the central and less central items in each task order for Experiment 1.

Task order	Centrality	
	Central	Less Central
<u>Centrality task</u>	4.72 (.69)	2.89 (.50)
<u>Preceding induction task</u>		
<u>Centrality task</u>	6.01 (.62)	3.94 (.54)
<u>Following induction task</u>		

Note. Higher estimates indicate higher centrality.

Two types of analyses were carried out - one treating participants as a random factor (F_1), the other items (F_2). Participants analyses allow generalization of the results to the participant, while the items analyses allow generalization to the item population (Clark, 1973). In the F_1 analysis Centrality was treated as a repeated factor, in the F_2 analysis Centrality and Task order were treated as repeated measures factors. Centrality had a main effect ($F_1(1,20) = 91.17, p < .001$; $F_2(1,11) = 557.73, p < .001$). The effect of Task order was not significant for participants ($F_1(1,20) = 2.06$) but it was for items ($F_2(1,11) = 197.33, p < .001$). The Centrality x Task order interaction was not significant for participants ($F_1 < 1$) but it was for items ($F_2(1,11) = 3.30, p < .05$).⁶

Induction task. For a particular participant, the member of each argument pair that had the highest centrality score was designated as "central", the other as "less central". The

⁶ The alpha level is set at .05 unless otherwise stated. ANOVA tables are presented in the Appendix.

scores from the forced-choice induction task were organized accordingly in the "central" and "less central" groups. Since choice of one member of an argument pair precluded the choice of the other, the two choices were not independent. Two types of analyses were carried out - one treating participants as a random factor the other items. For the participant analysis, the dependent variable was the proportion of times the central member was chosen in the forced choice induction task (chance = .50). For the item analyses, the dependent variable was the proportion of participants that chose the central member in the induction task (chance = .50).

As expected, arguments involving properties rated as more central were chosen more often in the induction task. For participants, the mean (SE) score was .62 (.039). For items, the mean (SE) score was .61 (.029).

Preliminary tests failed to show a main effect of task order (participants: $F_1(1,22) < 1$; items: $t(11) = -1.16$; two-tailed paired samples t-test)⁷. The data were therefore collapsed across task order. Properties rated as central were more projectible than properties rated as less central (participants: $t(23) = 3.14$, $p < .01$; items: $t(11) = 4.57$, $p < .01$; one-sample t-tests: chance = .50).

DISCUSSION

In Experiment 1 the relative centrality of pairs of properties was estimated from the ease-of-imagining task. The harder a participant found it to transform the image of an animal into one that lacked a specific property, the more conceptually central that property was deemed to be. The results showed that, as predicted, immutable properties are more projectible than mutable ones.

2.4 EXPERIMENT 2: MUTABILITY AND CENTRALITY IN A DEPENDENCY SPACE

One could object that the results of Experiment 1 are an artifact of the forced choice design. If participants made an independent decision about each argument, then possibly they

⁷ All t-tests reported in this thesis are two-tailed unless otherwise stated.

would be indifferent to projecting the central or the less-central property. Experiment 2 aimed to rule out this possibility by using the same materials as Experiment 1 while asking for conditional likelihood estimates for each argument in a pair separately. To the extent that conceptual centrality influences projectibility, we expected Experiment 2 to replicate the results of Experiment 1.

Experiment 2 used two tasks to estimate conceptual centrality: (i) the ease-of-imagining task, and (ii) the dependency structure task. To the extent that both measures gauge conceptual centrality we expected participants' responses to be independent of the task used.

METHOD

Participants. A new sample of 55 first-year psychology undergraduates from the University of Durham or from Northumbria University volunteered to participate.

Materials and Design

Induction task. The materials were borrowed from Experiment 1 (see Table 2.1). The only change was that in the present experiment participants were asked to provide conditional likelihood estimates for each argument in a pair. That is, the two arguments comprising an induction item of Experiment 1 were divided into two separate arguments. A sample item of the induction task was:

Assuming that giraffes have dense nerve fibres in the mid-brain, how likely do you think it is that zebras have dense nerve fibres in the mid-brain?

1	2	3	4	5	6	7	8	9
very likely				unsure				very unlikely

Centrality Estimation task. Two different measures were used to estimate the conceptual centrality of a feature, the ease-of-imagining task, and the dependency structure task. The materials for the ease-of-imagining task were exactly the same as those of

Experiment 1. The dependency structure task asked participants to evaluate how much about an animal depend on a given property. A sample item is presented below.

Assuming that giraffes have dense nerve fibres in the mid-brain, how much about giraffes do you think depends on this property?

1	2	3	4	5	6	7	8	9	
a great deal									almost nothing

All participants were first presented with the centrality estimation task followed by the induction task. Twenty-two participants had to estimate the centrality of the properties from the ease-of-imagining task (imagining condition), and 23 from the dependency structure task (dependency condition). The experiment was, therefore, a two factors repeated measures design. The order of the items in each task was separately randomized for each participant.

Procedure.

Induction task. A similar procedure was used to that of Experiment 1. Participants had to evaluate 24 items. They responded with probability judgments in a 9-point scale, representing the conditional likelihood that the conclusion category has the property in question given that the premise has. Lower estimates indicated higher judged probabilities. The instructions were similar to those of Experiment 1 with minimal changes to compensate for the fact that in the present study we asked for conditional likelihood estimates.

Centrality Estimation task. The imaging group followed the same procedure as the corresponding group in Experiment 1, except that the scale was reversed. In Experiment 2, a rating of "9" meant that participants found it very easy to imagine the animal lacking the property, whereas a rating of "1" meant that they found it very difficult to complete the transformation. Participants in the dependency group were informed that they would be presented with some assumptions about an animal followed by a question in which they were asked to estimate how much about an animal depended on a property in a 9-point scale. A rating of "1" meant that they thought that a great deal about the animal depended on the

property, whereas a rating of "9" meant that they believed that "almost nothing" about the animal depended on the property. The lower the estimate the more about an animal depended on the property. Participants were given an example to work through.

Unconditional probability task. Participants were instructed that they would see some statements whose likelihood they had to judge. This task used the same scale as the induction task.

RESULTS

Centrality Estimation task. Table 2.3 presents the mean centrality scores for the central and the less central items for each centrality estimation task. The mean centrality scores for central items are lower than the corresponding scores for the less central items. The mean scores for the Dependency task are lower than the corresponding scores for the Ease-of-imagining task.

Table 2.3 Mean (SE) centrality estimates for the central and less central items for each centrality measure for Experiment 2.

Centrality Measure	Centrality	
	Central	Less Central
<u>Ease-of-imagining</u>	4.46 (.41)	6.79 (.33)
<u>Dependency</u>	3.31 (.13)	5.61 (.20)

Note. Lower estimates indicate higher centrality.

Two 2(Centrality) x 2(Centrality measure) analyses of variance were carried out (one for F_1 and for F_2). In the F_1 analysis Centrality was treated as a repeated measures factor, in the F_2 both centrality and Task Order were treated as repeated measures factors. Centrality had a significant main effect ($F_1(1,43) = 263.83, p < .001$; $F_2(1,11) = 384.27, p < .001$). Centrality Measure had a significant main effect ($F_1(1,43) = 9.43, p < .05$; $F_2(1,11) = 27.19, p < .001$). No interaction was detected (both $F_s < 1$).

Induction and Unconditional Probability tasks. The likelihood scores from the induction and the unconditional probability tasks were organized in two distinct ways. Coding A was similar to the one used in Experiment 1. Items were divided into two groups according to their estimated centrality. Conditional and unconditional likelihood measures for the two conditions were compared. Coding B involved calculating for each participant a correlation coefficient between the centrality estimates and the conditional likelihood estimates while partialling out the effect of the unconditional likelihood estimates. Conditional likelihood estimates might be contaminated by prior knowledge. This coding examined whether conceptual centrality influences the strength of an argument beyond the effects of believability in the conclusion.

Coding A. Scores of the centrality estimation task were coded as in Experiment 1. For a given item, the property that was either harder to imagine an animal being without, or the property upon which more of an animal's functions depended, was coded as central while its counterpart as less central. The data from the induction and the unconditional probability tasks were organized in the resulting central and less central conditions.

Table 2.4 summarizes the results. On average, participants assigned both higher conditional and higher unconditional likelihood estimates for the central rather than the less-central properties.

Table 2. 4 Mean conditional (unconditional) likelihood estimates as a function of Centrality and Centrality Estimation task.

Centrality Measure	Centrality	
	Central	Less-Central
<u>Ease-of-imagining</u>	4.31 (4.35)	4.42 (4.69)
<u>Dependency</u>	4.03 (3.93)	4.52 (4.77)
	4.17 (4.13)	4.47 (4.73)

Note. Lower estimates indicate higher judged probabilities.

Induction task. As for the Centrality estimates, two 2 (Centrality) x 2 (Centrality Measure) analyses of variance were carried out. Centrality had a significant main effect ($F_1(1,43)=6.85, p<.05; F_2(1,11)=7.52, p<.05.$). Centrality Measure did not have a significant effect on either participants or items (both $F_s<1$). There was no significant interaction for participants ($F_1(1,43)=2.75$) but there was for items ($F_2(1,11)=5.62, p<.05$).

Unconditional probability task. Similar 2 (Centrality) x 2 (Centrality Measure) analyses of variance were performed on the unconditional likelihood estimates. Centrality had a significant main effect ($F_1(1,43)=13.57, p<.005; F_2(1,11)=17.43, p<.005$). No significant effect was observed for Centrality Measure ($F_1(1,43)=1.27; F_2(1,11)=3.71$). The interaction failed to reach significance ($F_1(1,43)=2.51; F_2(1,11)=3.94$).

The fact that centrality estimates are associated with unconditional likelihood estimates leaves open the possibility that the effect of centrality on conditional likelihood estimates is spurious. Participants seem familiar with some of the properties used. Possibly such familiarity influenced both their centrality and their conditional likelihood estimates. To examine this interpretation we re-coded the data.

Coding B. For each participant a correlation coefficient was computed relating the raw centrality estimates to their conditional likelihood counterparts while partialling out the effect of unconditional likelihood estimates. The mean (SE) was .17 (.034). Following Lorch and Myers' (1990) recommendation, a one-sample t-test was performed on the resulting coefficients to compare them to chance performance of 0; $t(44)=3.89, p<.001$. The raw data (for all experiments) are shown in the accompanying floppy disk.

DISCUSSION

Experiment 2 had two main purposes: (i) to detect whether the centrality effect would obtain with independent judgments, and (ii) to detect whether the ease-of-imagining and the dependency structure tasks measure the same thing. The results support both predictions. The more central a property the higher its judged inductive strength. This effect was significant even when the effect of unconditional likelihood estimates was partialled out (from both

factors). Therefore the centrality effect of Experiment 1 was not an artifact of the forced choice design. Furthermore, the results failed to show a differential effect of centrality depending on the centrality estimation task used. This supports the view that the ease-of-imagining and the dependency tasks gauge on the same factor.

2.5 GENERAL DISCUSSION

In Experiment 1 participants were presented with items consisting of a pair of categorical arguments which had the same premise and conclusion categories but involved different properties. For each item, participants were asked to choose the arguments that they found more convincing. Also, for each argument, participants were asked to estimate how easily they could imagine an animal lacking a property that it normally has. The results showed that participants prefer to project immutable over less immutable properties. Experiment 2 replicated the centrality effect while asking for conditional likelihood estimates separately for each argument that comprised an item of Experiment 1. Further, Experiment 2 used two distinct operational definitions of centrality: the ease-of-imagining definition and the dependency structure definition. Both definitions seem to tap the same underlying factor since the analyses showed no response difference depending on the definition used.

The following chapters manipulate the conceptual centrality of features directly by manipulating their dependencies. There are both theoretical and experimental reasons for doing this. Theoretically, immutability is not a well-articulated notion, centrality in a concept's dependency space is. Ultimately, a model of category-based inference can be constructed by building on the Sloman et al. (1998) model of concept centrality. Experimentally, manipulating the centrality of a (relatively unfamiliar) feature by manipulating its dependencies has the advantage of being direct; it allows us to manipulate rather than to estimate feature centrality. For all participants, a given feature would be (relatively) either central or less-central - no separate analyses of the data for each participant

are needed. Recall that the data of Experiments 1 and 2 were analyzed by comparing centrality estimates to inductive scores for each participant separately.

CHAPTER III:

CONCEPTUAL CENTRALITY FROM A DEPENDENCY CHAIN

Experiment 2 supported the view that the Sloman et al. (1998) notion of conceptual centrality as centrality in a concept's dependency space is a good surrogate of immutability. Henceforth, we manipulate relative feature centrality based on the Sloman et al. notion, i.e., we take a feature to be central to the extent that other (central) features depend on it. In the present chapter the relative centrality of two features is operationally defined by embedding them in a dependency chain. All else being equal, if feature C depends on feature A ($C \rightarrow A$), then feature A is relatively more central than feature C. For birds, *has wings*, for instance, is a more central feature than *can fly* because flying depends on wings, while wings do not depend on flying. Experiments 3, 4, and 5 define the relative centrality of two features by embedding them in such a dependency relation. Experiment 3 also provides concrete explanations for the dependency relations (e.g. exceptional hearing causes frequent headaches because very high auditory frequencies cause nasal cavities to vibrate). To the extent that the projectibility of a property depends on concrete explanations, we expected explanations to enhance the centrality effect. On similar lines, all else being equal, given that $C \rightarrow B$ and $B \rightarrow A$, feature A is relatively more central than feature C. Experiments 6 and 7 manipulate the relative centrality of two features by embedding them in such a dependency chain. Our straightforward prediction is that the depended-on properties (feature A in the schematic examples) will be more projectible than the dependent properties (feature C in the schematic examples) to the extent that the premise and conclusion categories share structure.

A reason to expect centrality effects with centrality defined from a single dependency chain comes from the domain of categorization. Ahn and Lassaline (1995; see also Ahn & Dennis, 1997) showed that causes are weighted more heavily than effects in categorization decisions (the *causal status hypothesis*). For instance, given a disease characterized by symptoms A, B, and C such that A causes B which causes C, participants were more likely to diagnose a person with the disease given symptoms A and B but not C, rather than given symptoms B and C but not A. Causes are more central than effects because effects depend on causes, whereas causes do not necessarily depend on effects (though they may provide

evidence for them). These studies can therefore be said to show that central properties are weighted more strongly than less-central ones in categorization decisions.

As we already mentioned in Chapter II, inference is a major function of categorization. Having categorized an object as a member of a category sanctions lots of inferences about the object. Further, knowing that an object has a novel property leads us to infer that other similar objects will have the property as well. To the extent therefore that centrality influences categorization, we also expect centrality to affect inter-categorical inference.

The present experiments investigate four broad issues:

1. The role of centrality on projectibility
2. The role of shared dependency structure: Does it mediate centrality effects?
3. The role of vagueness: Does centrality operate under vague conditions?
4. Feature centrality: centrality in a causal or a dependency space?

Below we unpack issues 2, 3, and 4 in detail.

3.1 SHARED DEPENDENCY STRUCTURE

The introductory chapter stated that the role of centrality on projectibility should be mediated by the extent to which the premise and conclusion categories share structure. This claim is based on the observation that centrality is concept-relative, one and the same feature may have different centrality statuses across concepts (e.g., Medin & Shoben, 1988). The fact that centrality is concept-relative can be explained by our appeal to dependency structure. Different concepts involve features embedded in different dependency networks, and as a result the centrality status of a feature across concepts may differ. The color of a solar panel is one of its central properties, the color of a soccer-ball is not.

Although we expected the effect of centrality to be mediated by the extent to which the premise and conclusion categories shared structure, we had no idea about how much structure the categories should share. This issue was addressed empirically by using premise-

conclusion category pairs that are associated with different similarity estimates. We assumed that the more similar a pair of animals was judged, the more dependency structure their concepts shared. Experiments 3 and 4 manipulated common structure by using arguments involving either high or low similarity mammal pairs. Experiments 6 and 7 manipulated shared structure more widely. They used animal pairs high in similarity and from the same superordinate (HSSS; e.g. lions - tigers), low in similarity from the same superordinate (LSSS; e.g., elephants - mice), or low in similarity from different superordinates (LSDS; e.g., tigers - falcons).

Experiments 6 and 7 also addressed two claims about the role of taxonomies in property induction: (1) that sharing a superordinate decides the strength of arguments, and (2) that sharing a superordinate mediates the centrality effect on property induction. If claim (1) were true, then we should not observe differences in property generalizability among HSSS and LSSS categories, but we should observe differences among HSSS or LSSS and LSDS categories. Similarly, if (2) were true, then we should not observe differences in the generalizability of the central and the less central properties among same superordinate categories, but we should observe differences among different superordinate categories.

3.2 VAGUENESS AND DIRECTIONALITY

In the introductory chapter we discussed structural-mapping models of category-based inference, models that view object comparison as a process of alignment between structured representations. Properties (and relations) connected to the resulting aligned structure which are present in the base but absent in the target category constitute potential inferences. According to these models, property induction is constrained by the principle of systematicity. All else equal, the higher the order of relation that a candidate property is embedded in, the higher its projectibility. The principle of systematicity as an inductive constraint has gained empirical support in recent studies (Lassaline, 1996; Wu & Genter, 1998).

In a sense, evidence that conceptual centrality affects projectibility, is also evidence for the principle of systematicity as an inductive constraint. Properties that are more central in the dependency structure of a concept are generally those more deeply embedded in the hierarchical structure of relations that structural mapping models' hold concepts consist of. A critical difference is our appeal on the status of a property in a relation. The directionality of relations are critical in determining centrality; in the simple $B \rightarrow A$ case, all else equal, feature A to be more central than feature B. In contrast, the principle of systematicity does not consider the status of a property in a relation. For the $B \rightarrow A$ case, for instance, it does not predict that feature A has more inductive potential than feature B. The experiments reported in this chapter directly assess this hypothesis.

Finally, notice that for structural-mapping models to work, predicates, as well as their relations to other predicates, need all to be completely specified. Otherwise stated, such models cannot deal with vagueness. Experiment 5 investigated centrality effects with abstract properties, i.e., under conditions of vagueness.

3.3 FEATURE CENTRALITY: CENTRALITY IN A CAUSAL OR A DEPENDENCY SPACE?

Ahn and Lassaline (1995) have made a similar proposal about feature centrality, what they call the *causal status hypothesis* (see also Ahn, 1998). According to this hypothesis, causal features are more central than effect features in category decisions. Following Sloman et al. (1998), we do not give a special status to causal over other asymmetric dependency relations as determinants of feature centrality. To avoid misunderstanding, possibly causal relations are in fact stronger (in the Sloman et al. model this translates into higher d_{ij} values). However, we hold that all asymmetric dependency relations are similar in how they contribute towards a feature's centrality. Experiment 5 examines this hypothesis by defining relative feature centrality from a generic dependency chain, e.g., "property W depends on property Z".

3.4 EXPERIMENT 3: CENTRALITY FROM A CAUSAL + EXPLANATION RELATION

Experiment 3 aimed to test two hypotheses: (1) to detect whether a local causal relation together with an explanation suffices for centrality effects, and (2) in the case it does, it aimed to detect a possible boundary condition for the centrality effect set by similarity. To test hypothesis (1) 16 categorical arguments were constructed such that for half of them a causal relation was established and an explanation for this relation was given (the causal+explanation condition), whereas for the rest no relation was stipulated (the control condition). Hypothesis (2) was addressed empirically by systematically varying the similarity between the premise and conclusion animal pairs. Half of the items involved pairs of mammals that were highly similar, whereas the other half included pairs that had low similarity. We expected that the participants would be more willing to project central over less-central properties (causes over effects in the causal+explanation condition) rather than the same properties in the control condition.

METHOD

Participants. Participants were 24 University of Durham students.

Design and Materials. Relation (causal+explanation or control) was crossed with Shared structure (high similarity or low similarity mammals) in a within subjects design on both factors.

Participants were given 16 items to evaluate. Half of the items involved high similarity mammal pairs (e.g. horse-cow), whereas the other half involved low similarity mammal pairs (e.g. dolphin-cow). Eight mammal pairs were used, each of which appeared in 2 orders (e.g. horse-cow, cow-horse). See top of Table 3.1 for the complete set of premise-conclusion categories.

Each of the resulting 16 mammal-pairs was assigned a pair of properties (for examples see Table 3.2; the full set of properties used is shown in Table A.3.1). At each level

of shared structure, for half of the items a causal relation was stipulated and an explanation for this relation was given (the causal+explanation condition), whereas for the rest no relation was stipulated (the control condition). At each level of shared structure, the same 4 mammal pairs were used to form the causal+explanation and control conditions except that their assignment to premise and conclusion categories was swapped across conditions. Four lists of materials were constructed to counterbalance the assignment of mammal-pairs to the causal+explanation and control conditions (see Table 3.1).

Table 3. 1 Design and Materials of Experiments 3 and 4. Experimental = causal+explanation condition (Exp. 3), causal condition (Exp. 4).

	High Similarity		Low Similarity	
	Cow-Horse	Horse-Cow	Seal-Horse	Horse-Seal
	Squirrel-Mouse	Mouse-Squirrel	Elephant-Mouse	Mouse-Elephant
	Seal-Dolphin	Dolphin-Seal	Cow-Dolphin	Dolphin-Cow
	Elephant-Rhino	Rhino-Elephant	Squirrel-Rhino	Rhino-Squirrel
List 1	Experimental	Control	Experimental	Control
List 2	Experimental	Control	Control	Experimental
List 3	Control	Experimental	Experimental	Control
List 4	Control	Experimental	Control	Experimental

For each argument, participants were asked to choose the property (cause or effect in the causal condition) that the target mammal was more likely to have. Table 3.2 below illustrates samples from each Shared structure by Relation condition.

Both the position of the properties after "Fact" and their evaluation order were separately randomized for each participant. Table 3.2 shows only one position of properties after "Fact" (the one where the central property is presented first), and only one evaluation order (the one where the central property appears as the first choice).

Table 3. 2 Sample items from each Shared Structure by Relation condition of Experiments 3 and 4. Control items contained the description in bold- faced letters. Causal+explanation items (Exp. 3) contained the whole description. Causal items (Exp. 4) contained the whole description except the text inside brackets.

Shared structure

High Similarity

Fact: Squirrels have exceptional hearing and frequent headaches.

For squirrels, exceptional hearing causes frequent headaches [because very high auditory frequencies cause nasal cavities to vibrate.]

Choose the statement below that you consider more likely to be true.

A. Mice have exceptional hearing.

Central

B. Mice have frequent headaches.

Less central

Low similarity

Fact: Squirrels have an underdeveloped hippocampus and anterograde amnesia.

For squirrels, an underdeveloped hippocampus causes anterograde amnesia [because underdeveloped mammillary bodies adversely affect memory.]

Choose the statement below that you consider more likely to be true.

A. Rhinos have an underdeveloped hippocampus.

Central

B. Rhinos have anterograde amnesia.

Less central

Procedure. Participants were presented with a booklet, containing one of the resulting 4 lists of material. There was no time limit but participants were encouraged to work quickly. Participants read the following instructions:

Imagine that you are a first-year undergraduate in biology interested in the physiology and behaviour of mammals. You recently got hold of an authoritative book on the biology of mammals only to find out that some of the pages were missing or were torn apart. You are left with some excerpts from the book stating facts about the biology of some mammals. You are informed that a given mammal has two properties. These properties may refer to any physiologically meaningful aspect of that mammal. Beneath each excerpt you are presented with a different mammal and you are asked to choose the property that this mammal is more likely to have.

Participants had to work through two examples similar to the test items, one in the causal+explanation and another in the control condition. They were also asked to treat each test example separately.

RESULTS

Coding. The choice of the central property (the cause in the causal+explanation condition) was coded as 1, and the choice of the less central property (the effect in the causal+explanation condition) as 0. The control items were scored accordingly: control items whose causal+explanation counterpart was central were coded as 1, otherwise as 0. The dependent variable was the proportion of central property choices over the total number of choices. Each participant contributed 4 data points, one in each Shared structure by Relation condition.

Table 3.3 summarizes the results. Overall, participants chose more the central over the less-central properties in the causal+explanation condition, rather than the corresponding properties in the control condition. Further, the preference to project central over less-central properties is higher for arguments involving high similarity rather than low similarity categories.

Table 3.3 Mean proportion choice of the property designed as the cause in the causal+explanation condition.

Relation	Shared Structure	
	High Similarity	Low Similarity
<u>Causal</u>	.70 (.06)	.61 (.05)
<u>Control</u>	.56 (.06)	.53 (.04)

Preliminary tests showed that the different lists did not have an influence on participants' responses. The data were therefore collapsed across lists.

Two 2 (Shared structure) x 2 (Relation) analyses of variance were carried on the

resulting data (one for F_1 and one for F_2). In the F_1 analysis both factors were repeated measures, in the F_2 analysis only Relation was a repeated measures factor. Shared structure failed to reach significance ($F_1(1, 23) = 1.21$; $F_2 < 1$). Relation had a robust main effect ($F_1(1, 23) = 4.29$, $p = .05$; $F_2(1, 14) = 13.24$, $p < .005$). There was no significant Shared structure x Relation interaction (both $F_s < 1$).

DISCUSSION

In sum, only the centrality manipulation had a significant effect. Participants were more willing to project properties designed as causes rather than effects in the causal+explanation condition, rather than their counterparts in the control condition. The fact that the centrality effect was stronger in the item analyses is unsurprising. Participants could possibly reason about some of the properties used, and hence believability and variability effects came into play. Recall that the participant analysis compared responses for different items. In the item analysis, where each causal+explanation item was pitted directly against its control counterpart, all extraneous effects are cancelled out. That is, the item analysis provides a cleaner test for our hypothesis.

The significant effect of centrality may be attributable to the causal relation, or also to the explanation provided. Experiment 4 directly assessed this hypothesis by using the same items as Experiment 3, while dropping the explanations from the causal+explanation condition to form a causal condition. To the extent that supplying explanations play a role, we expected participants to be less willing to project the central property in the causal condition. If the opposite is the case, we expected no differences in the results of the two experiments.

The fact that shared structure failed to reach significance, suggests that centrality effects defined from a single causal+explanation relationship are not mediated by similarity. Possibly we failed to detect such an interaction because all category pairs were relatively similar; they all involved mammal categories. Experiments 6 and 7 manipulated shared structure more widely by manipulating superordinate membership.

3.5 EXPERIMENT 4: CENTRALITY FROM A CAUSAL RELATION

The results of Experiment 3 show a robust effect of centrality. Experiment 4 aimed to investigate whether this effect was due to the causal relationship, or also to the explanation given. To this end, Experiment 4 used the same items and procedure as Experiment 3 with the change that the explanations were dropped from the causal+explanation condition to form what we call the causal condition.

METHOD

Participants. A new sample of 24 University of Durham students voluntarily participated in this experiment. The experiment was run in the computer room of Durham's Main Library.

Design and Materials. Experiment 4 used the same design as Experiment 3. Relation (causal or control) was crossed with Shared structure (high similarity or low similarity mammals) in a within subjects design on both factors.

The materials of this experiment were similar to those used in Experiment 3, except that the explanations were dropped from the causal+explanation condition and formed the causal condition. The control conditions were identical across experiments. For sample items in each Shared structure by Relation condition refer back to Table 3.2. (For the complete list of items see Table A.3.1)

Procedure. The procedure of Experiment 4 used was similar to that of Experiment 3, with the change that the instructions of Experiment 3 were minimally revised to compensate for the change from the causal+explanation to the causal condition. The causal+explanation example in the instructions gave way to its causal counterpart.

RESULTS

Coding. The data were coded in the same way as in Experiment 3. The choice of the central property (the cause in the causal condition) was coded as 1, and the choice of the less central property (the effect in the causal condition) as 0. The dependent variable was the ratio of central property choices over total number of choices. Each participant contributed 4 data points, one in each Shared structure by Relation condition.

Table 3.4 summarizes the results. Overall, participants preferred to project the central over the less-central properties in the causal condition, rather than the corresponding properties in the control condition. Further, the preference to project central over less-central properties is higher for arguments involving high similarity rather than low similarity categories.

Table 3. 4 Mean (SE) proportions of the property designed as the cause in the Causal condition.

Relation	Shared Structure	
	High Similarity	Low Similarity
<u>Causal</u>	.77 (.06)	.66 (.05)
<u>Control</u>	.63 (.05)	.51 (.05)

Preliminary tests showed that the different lists used did not affect participants' responses. The results were therefore collapsed across lists.

As in Experiment 3, two 2 (Shared structure) x 2 (Relation) analyses of variance were carried out. Shared structure had a main effect for participants ($F_1(1, 23) = 8.62, p < .01$), but not for items ($F_2(1, 14) = 1.74, p > .20$). Relation had a robust main effect ($F_1(1, 23) = 6.19, p < .05$; $F_2(1, 14) = 15.96, p < .001$). No significant interaction was found (both $F_s < 1$).

DISCUSSION

The results show that centrality had a robust effect; i.e. participants were more willing to project central properties in the causal condition rather than these properties' counterparts in the control condition. Shared structure did not have a robust effect. Furthermore, since the means in the causal condition are higher than those in the causal+explanation condition of Experiment 3, we can conclude that the results are indeed produced by the causal relation and not by the explanation given.

It may be the case that the explanations given in Experiment 3 were somehow unbelievable or convoluted and hence they failed to further strengthen the projectibility of the central property. Furthermore, it is possible that the participants in the latter experiment created their own explanations to account for the causal relations. One's own explanations may have a bigger effect on judgment than explanation by others. In any case, stipulating a causal relation is sufficient to produce centrality effects.

3.6 **EXPERIMENT 5: CENTRALITY-IN-THE-ABSTRACT**

Experiment 5 has two main objectives: (1) to detect whether a dependency relation suffices for centrality effects, and (2) to investigate whether centrality operates under conditions of vagueness.

METHOD

Participants. A new sample of 32 University of Durham students volunteered to participate.

Materials, Procedure, and Design. Participants were asked to make a choice for the following argument that appeared in a single page:

Fact: Mammal A has both physiological properties Z and W.

For mammal A, property W depends on property Z.

Choose the statement below that you consider more likely to be true.

1. Mammal B has Z.
 2. Mammal B has W.
-

No instructions were given, but we assume that the participants treated mammals A and B as different. Both the presentation order of the properties in the first statement (Fact:...) and the two response alternatives were counterbalanced. Z always referred to the central property and W to the less-central one. The experiment used a goodness-of-fit χ^2 design.

RESULTS AND DISCUSSION

One participant wrote “both the same” and hence was excluded from the analysis. Out of the remaining 31 participants, 22 chose to project the central property (Z). The binomial probability of getting at least 22 out of 31 Z choices by chance is less than .02 ($\chi^2(1) = 5.45$, $p < .02$).

Centrality effect. The results suggest that stipulating a single dependency relation is sufficient for centrality effects. It seems therefore that centrality in a dependency space (as we propose) rather than centrality in a causal space (as theorists of the causal status hypothesis propose) better captures the results.

Shared structure effect. Assuming that participants considered mammals A and B as different, once again it seems that similarity does not affect the willingness to project the central (depended-on) versus the less-central (dependent) properties. Interestingly, the magnitude of the centrality effect (the number of central over less-central choices) is closer to those of the high-similar conditions of Experiments 3, and 4. Possibly, in the absence of any information about the species involved, participants considered both mammals as similar since the label “mammal” was particularly salient. Empirical evidence suggests that the use of common labels facilitates category-based inference (Gelman & Markman, 1986; Gelman, 1989; Davidson & Gelman, 1990). For instance, Gelman and Markman (1986) have shown that giving two creatures a common linguistic label (like “bird”) increased preschoolers propensity to import knowledge from one creature to another.

Vagueness. In the current experiment both categories and properties were left abstract. The results support the view that feature centrality operates even under vague conditions. It is difficult to see how structural mapping models of category-based induction

can account for such results, since for these models to work need a complete specification of the predicates used, as well as their relations to other predicates.

Taken together, the results of experiments 3, 4, and 5 show that central properties are more projectible than less-central ones in category-based inferences involving mammals. This effect was independent of shared structure. This may be because shared structure was not widely varied (both conditions involved mammals). A number of objections can be raised against our interpretations of the results. One objection is that our results might be due to a demand characteristic because they stem from a forced choice design. It is possible that, given the option, people will be indifferent about projecting the central or the less-central properties. (Recall that Experiment 2 suggested otherwise). A further objection is that the central properties came from different general classes than the less-central ones (Experiments 3 and 4). Recall though that our design controlled for this discrepancy by including the control condition.

3.7 EXPERIMENT 6: CENTRALITY FROM A SINGLE DEPENDENCY CHAIN

Experiment 6 aimed to control for possible objections to the previous experiments. First, participants were asked to provide independent likelihood estimates for the central and the less-central properties. Second, both properties for a specific item came from the same general class (e.g., both were either enzymes, hormones, or neurotransmitters). Further, the range of shared structure was extended - we tested animals from different superordinates (e.g., falcons-hippos) as well as from the same superordinate (e.g. rhinos-hippos). Finally, we used a wider variety of dependency relations (e.g., depends, regulates).

METHOD

Participants. A new sample of 33 University of Durham undergraduates voluntarily participated in the study.

Design and Materials. Experiment 6 crossed Centrality (central or less-central properties) with Shared structure (high-similar same-superordinate (HSSS), or low-similar same-superordinate (LSSS), or low-similar different superordinate (LSDS) animal pairs) in a repeated measures design on both factors.

Animal pairs of different shared structure were selected from the Osherson et al. (1990) study where they provide a pair-wise similarity matrix (p. 196). However, for some items we relied on our intuitions. A check on the similarity of animal pairs was carried out by an independent group of 11 University of Durham students. The students were asked to rate the biological similarity of the 18 animal pairs of Experiment 6 on a 0 to 10 scale, where 0 was labeled as “highly dissimilar” and 10 as “highly similar”. The mean similarity ratings were 7.86 for the HSSS condition, 3.26 for the LSSS condition, and 1.75 for the LSDS condition. The animal pairs used are presented Table 3.5 as a function of decreasing similarity. Notice that similarity decreases as one moves from HSSS, to LSSS, to LSDS animal pairs. The similarity estimates mark well the borders of the similarity conditions, with the exception of the Deer-Dolphin pair of the LSSS condition that overlaps in similarity with pairs of the LSDS condition. It seems that sharing a common superordinate is not sufficient to discriminate similarity estimates, below we argue that it is neither critical in projecting a property amongst categories or in mediating centrality effects.

Table 3. 5 Animal pairs of Experiment 10 along with the mean (SE) similarity ratings.

HSSS	SIM	LSSS	SIM	LSDS	SIM
Gorilla-Chimp	9.2 (.30)	Beaver-Hippo	3.9 (.77)	Sparrow-Tiger	2.2(.63)
Lion-Tiger	8.5 (.39)	Raccoon-Tiger	3.7 (.76)	Blackbird-Chimp	2.1(.69)
Rhino-Beaver	8.2 (.30)	Zebra-Chimp	3.5 (.86)	Robin-Horse	1.8(.58)
Seal-Dolphin	7.2 (.52)	Ferret-Horse	3.2 (.81)	Eagle-Mouse	1.7(.60)
Cow-Horse	7.2 (.81)	Bear-Mouse	3.1 (.76)	Falcon-Hippo	1.5(.59)
Squirrel-Mouse	6.9 (.56)	Deer-Dolphin	2.2 (.68)	Swallow-Dolphin	1.2(.44)

Participants were informed that an animal had two properties: one upon which one of its functions depended (the central property), and one that depended upon that function (less central). Participants were asked to estimate the likelihood that another animal had each of these properties. Eighteen items were constructed, 6 in each shared structure condition. Table 3.6 presents sample items from each Shared structure by Centrality conditions. For the full list of properties used see Table A.3.6.

Table 3. 6 Sample items from each Shared Structure by Centrality condition of Experiment 6. Only one name order and one premise triple by property type assignment is shown.

Shared structure

HSSS Fact: Rhinos have both the enzyme lipase and the enzyme amylase.
 For rhinos, the enzyme lipase regulates metabolism.
 For rhinos, the enzyme amylase is regulated by metabolic rate.

Rate the likelihood of the following statements.

- A. Hippos have lipase. Central
- B. Hippos have amylase. Less central

LSSS Fact: Beavers have both the hormone prolactine and the hormone renin.
 For beavers, the hormone prolactine regulates blood flow.
 For rhinos, the hormone renin is regulated by blood flow.

Rate the likelihood of the following statements.

- A. Hippos have prolactine. Central
- B. Hippos have renin. Less central

LSDS Fact: Falcons have both the neurotransmitter acetylcholine and the neurotransmitter noradrenalin.
 For falcons, the neurotransmitter acetylcholine helps detect predators.
 For falcons, the levels of the neurotransmitter noradrenalin increase after seeing a predator.

Rate the likelihood of the following statements.

- A. Hippos have acetylcholine. Central
- B. Hippos have noradrenalin. Less central

Note. HSSS=high similarity same superordinate. LSSS=low similarity same superordinate. LSDS=low similarity different superordinate.

Items were constructed as follows. The 18 premise categories were organized into 6 sets of triples, each triple consisting of an item from each of the three shared structure

conditions. For example, the three items shown in Table 3.6 comprise one triple. All three categories of a premise triple for were paired with a conclusion category to form one item in the HS-SS, one in the LS-SS, and one in the LS-DS shared structure condition, respectively. Table 3.7 shows the construction of each of the six triples.

Table 3.7 The 18 premise-conclusion categories of Experiments 6, 7, 10, 11 (first 3 triples), 15, and 16.

Triple #	Premise Triples			Conclusion
	<u>HSSS</u>	<u>LSSS</u>	<u>LSDS</u>	
1	Rhino	Beaver	Falcon	Hippo
2	Squirrel	Bear	Eagle	Mouse
3	Cow	Ferret	Robin	Horse
4	Gorilla	Zebra	Blackbird	Chimp
5	Lion	Raccoon	Sparrow	Tiger
6	Seal	Deer	Swallow	Dolphin

Table 3.8 Design of the 3 lists of materials of Experiments 6, 7, and 10.

Triple #	List1	List2	List3
1	EHN	HNE	NEH
2	EHN	HNE	NEH
3	HNE	NEH	EHN
4	HNE	NEH	EHN
5	NEH	EHN	HNE
6	NEH	EHN	HNE

Note. EHN = HSSS enzyme, LSSS hormone, LSDS neurotransmitter.
HNE = HSSS hormone, LSSS neurotransmitter, LSDS enzyme.
NEH = HSSS neurotransmitter, LSSS enzyme, LSDS hormone.

Three types of properties were used: enzymes, hormones, and neurotransmitters. For a given item, both properties came from the same biological class (e.g. both were enzymes). Arguments constructed from a given triple were assigned different property types. For example, for triple number 1, one member of the triple would have enzyme properties,

another would have hormone properties, and the third neurotransmitters. Then, a counterbalancing procedure made sure that each member of the triple occurred equally often with each property type to form three lists of materials. Table 3.8 above shows the design of the 3 lists of materials. Each of these lists had a counterpart where the names of the central and the less-central properties were reversed. Table 3.9 shows the properties used for each triple.

Table 3.9 Properties used in Experiments 6 and 7.

Triple #	Property Type		
	Enzymes	Hormones	Neurotransmitters
1	Lipase; protease	Prolactin; renin	Acetylcholine; nonadrenaline
2	Amylase; streptokinase	ACTH; LH	GABA; glycine
3	Papain; trypsin	TSH; FSH	Octopamine; histamine
4	Thrombin; subtilisin	ADH; MSH	Norepinephrine; epinephrine
5	Elastase; aliesterase	Oxytocin; calcitocin	Tyrosine; taurine
6	Phosphorylase; transacetylase	Vasopressin; proctolin	Quisqualate; kainate

Procedure. The procedure of Experiment 6 were similar to that used in Experiment 3, with the exception that the instructions of Experiment 3 were minimally revised to compensate for the fact that the present experiment involved animals. Participants worked through an example before starting the experiment.

RESULTS

The results are summarized in Table 3.10. As predicted, central properties were more projectible than less-central ones though this difference is small in absolute terms.

Table 3. 10 Mean (SE) inductive strength (mean likelihood estimates) for the Shared structure by Centrality conditions.

Shared Structure	Centrality	
	Central property	Less-Central property
<u>Same-Superordinate High-Similarity</u>	74.9 (2.11)	73.5 (2.12)
<u>Same-Superordinate Low-Similarity</u>	55.2 (2.99)	52.8 (3.01)
<u>Differ.-Superordinate Low-Similarity</u>	51.8 (2.92)	50.1 (2.98)

Two 2 (Centrality) x 3 (Shared structure) analyses of variance were carried out. In the F_1 analysis both factors were repeated measures, in the F_2 analysis only Centrality was a repeated measures factor. The analyses showed a main effect of Centrality ($F_1^*(1,32)=4.74$, $p<.05$; $F_2(1,15)=12.42$, $p<.005$) and a main effect of Shared structure ($F_1^*(2,31)=30.91$, $p<.001$; $F_2(2,15)=27.73$, $p<.001$)⁸. There was no significant interaction (both $F_s<1$).

To examine the locus of the main effect of shared structure, unplanned multiple pairwise comparisons were carried out using Tukey's HSD test. Across both participants and items the mean for the HSSS condition was significantly higher than that for either the LSSS or LSDS conditions. The means for the LSSS and LSDS conditions were not significantly different from each other⁹.

DISCUSSION

The main purpose of the present experiment was to control for objections that could be raised for our previous experiments. It asked for likelihood estimates, it matched the general class of the central and less-central properties, it varied the type of relation used, and it used a wider range of similarity. The results replicated the findings of the previous studies. Central properties were slightly, though significantly, more projectible than less-central ones. Shared

⁸ Whenever a repeated measures factor has three or more levels we report the exact F statistic (F^*) from the conservative Pillai's trace test (see Howell, 1997, p. 498).

structure had a significant main effect that was due to the differences between the HSSS condition and the LSSS or LSDS conditions. Again no centrality by shared structure interaction was evidenced. However, such an interaction should be obtained by using highly dissimilar animal pairs. Something that is central for lions, for instance, should not be projectible to insects, since these categories involve very different mental representations.

The fact that the effect of centrality was small might be due to background knowledge effects. Keeping the general class of properties within each item constant might have adversely stipulated relations that contradict general knowledge. For example, although we stated that an enzyme is controlled by metabolic rate, participants may still have thought that the enzyme also regulated metabolism thus breaking down the intended asymmetry.

On a different note, one could still object that the results of this experiment were due to a demand characteristic. The central and less-central properties of each item were estimated side-by-side, and this might have caused participants to focus on the dependency status of a property that would not have otherwise influenced their judgments.

3.8 EXPERIMENT 7: REPLICATION OF EXPERIMENT 6

The major aim of Experiment 7 was to detect whether the centrality effect of Experiment 6 is robust. Experiment 7 further addressed 2 issues: (1) whether the centrality effect of Experiment 6 was due to a demand characteristic, and (2) whether the weak centrality effect was due to participant's misinterpreting asymmetric relations as bi-directional ones. To address issue (1), each item of Experiment 6 was split to form two separate items - one about the central and one about the less-central property (this resembles our manipulation in Experiment 2). Issue (2) was addressed by instructing participants that the relations should be taken as unidirectional; e.g., when an item informed that an enzyme was regulated by metabolism that did not mean that metabolism was regulated by that enzyme.

⁹ For items, the Tukey's tests were carried out using SPSS. For participants, the tests were carried out by the author (mean HSSS = 74.21; mean LSSS = 53.97; mean LSDS = 50.98; CD=6.05 [MS_{error}=104.40; n=33; q_r(3, 64)=3.40]).

METHOD

Participants. A new sample of 33 University of Durham students volunteered to participate.

Materials and Design. Experiment 7 used the same design as that used in Experiment 6. Experiment 7 also used similar materials as those used in Experiment 6, except that each item of Experiment 6 was judged independently. For example, the HSSS item of Table 3.6 gave rise to the following 2 separate items:

A central item

Fact: For rhinos, the enzyme lipase regulates metabolism.

Rate the likelihood of the following statement:

Hippos have lipase. _____ %

Its less-central counterpart

Fact: For rhinos, the enzyme amylase is regulated by metabolic rate.

Rate the likelihood of the following statement:

Hippos have amylase. _____ %

The central and less-central counterparts of a single item of Experiment 6 were placed at maximum distance from one another (18 items apart).

RESULTS AND DISCUSSION

The results are summarized in Table 3.11. The main trends of Experiment 7 are very similar to those of Experiment 6. Centrality had a weak absolute effect (though higher than that of Experiment 6), shared structure had a significant effect, and there does not seem to be a Centrality by Shared structure interaction.

Table 3. 11 Mean (SE) inductive strength (mean likelihood estimates) for the Shared structure by Centrality conditions. Super=Superordinate.

Shared Structure	Centrality	
	Central property	Less-Central property
<u>Same-Super High-Similarity</u>	72.3 (2.87)	67.1 (3.56)
<u>Same-Super Low-Similarity</u>	47.5 (4.05)	44.1 (3.97)
<u>Differ.-Super Low-Similarity</u>	43.0 (3.74)	37.9 (3.60)

As in Experiment 6, two 2 (Centrality) x 3 (Shared structure) analyses of variance were carried out. The analyses showed a main effect of Centrality ($F_{1(1,32)}=4.75$, $p<.05$; $F_{2(1,15)}=19.55$, $p<.001$) and a main effect of Shared structure ($F_{1^*(2,31)}=46.50$, $p<.001$; $F_{2(2,15)}=42.16$, $p<.001$). There was no significant interaction (both $F_s<1$).

To examine the locus of the main effect of shared structure, unplanned multiple pairwise comparisons were carried out using Tukey's HSD test. Across both participants and items the mean for the HSSS condition was significantly higher than that for either the LSSS or LSDS conditions. The means for the LSSS and LSDS conditions were not significantly different from each other (for participants: mean HSSS=69.68; mean LSSS=45.78; mean LSDS=40.47; $CD=6.39$ [$MS_{error}=116.60$; $n=33$; $q_r(3, 64)=3.40$]).

The main purpose of Experiment 7 was (1) to replicate the centrality effect of Experiment 6, (2) to test whether this effect was due to a demand characteristic, and (3) to test whether the weak effect of centrality was due to participants misinterpreting the stipulated relations as bi-directional ones. A quick glance at the cells of Tables 3.9 and 3.10 shows that the trends of both studies are identical. In sum, the main effect of centrality was replicated and the results tentatively rule out interpretations (2) and (3). Also, the main effect of shared structure is due to differences between the HSSS and the LSSS or LSDS conditions. The small effect of centrality in both experiments might be in part because participants read the passive versions (statements concerning some of the less central properties) as actives. Further experimentation is needed to examine this hypothesis.

3.9 GENERAL DISCUSSION

Experiments of Chapter III addressed four broad issues: (1) the effect of feature centrality on projectibility, (2) the influence of shared structure between the premise and conclusion categories on centrality effects, (3) the role of vagueness and relational status of a feature on projectibility, and (4) whether feature centrality is best captured as centrality in a dependency or a causal space. Below we relate the current findings to each issue in turn.

3.9.1 Centrality effect

Throughout the experiments of Chapter III, centrality was defined in terms of a dependency relation. In each and every experiment participants were more willing to project the central over the less-central properties.

In Experiments 3 and 4 centrality was defined from a cause-and-effect relation; the cause being the central and the effect the less-central property. As predicted, causes were more projectible than effects. In Experiment 5 centrality was defined from the vague dependency relation "substance W depends on substance Z". The depended-on property (Z) was the central and the dependent (W) the less central property. The results showed that the central property (Z) was more generalizable than the less-central one (W). Finally, in Experiments 6 and 7 centrality was defined in terms of various types of dependency relations. Once again feature centrality had a strong effect on people's willingness to project a property.

3.9.2 Shared structure

The effect of feature centrality on projectibility is one aspect of our hypothesis. A second aspect is that this effect should be mediated by the extent to which the premise and conclusion categories share structure. To this end, we varied the similarity between the premise and conclusion categories. The more similar two categories were judged, the more structure they were assumed to share. In Experiments 3 and 4 some arguments involved similar while others dissimilar pairs of mammals. The effect of centrality on projectibility was independent from the level of shared structure. We reasoned that the failure to detect an

interaction might be because the amount of shared structure was not widely varied - all premise-conclusion pairs involved mammal categories. Experiments 6 and 7 examined this hypothesis by varying shared structure more widely. Some category pairs were highly-similar from the same superordinate, some low-similar from the same superordinate, and some low-similar from different superordinates. Despite this wider manipulation of shared structure, we once again failed to detect an interaction between shared structure and centrality on people's willingness to project properties.

Two factors may contribute to our failure to observe such an interaction: (1) that similarity was still not widely varied - all premise-conclusion pairs involved animals, and or (2) that the effect of centrality might depend on a stricter sense of similarity - on whether the conclusion category has the properties specified for the candidate feature for the premise. Notice that all the properties that depended upon the candidate features were specific and fairly general (e.g., metabolism, or heart rate). Hence, all animals could be safely assumed to share these properties. In Chapter IV we examine hypothesis (2) by using abstract dependent properties such as "lots of functions" or "few functions". If people project a property to the extent they believe that the conclusion category shares the properties that depend-on the candidate inference, then in this case we expect shared structure to interact with centrality.

As mentioned in the introduction of the present chapter, Experiments 6 and 7 by manipulating shared structure addressed two claims concerning the role of taxonomies in property induction: (1) that sharing a superordinate decides the strength of arguments, and (2) that sharing a superordinate mediates the centrality effect on property induction. If claim (1) were true, then we should not observe differences in property generalizability among HSSS and LSSS categories, but we should observe differences among HSSS or LSSS and LSDS categories. Contrary to claim (1), the results showed significant differences between the HSSS and the LSSS or LSDS conditions only. If claim (2) were true, then we should not observe differences in the generalizability of the central and the less central properties among same superordinate categories, but we should observe differences among different superordinate categories. Contrary to claim (2), no differences were detected (no shared

structure by centrality interaction). It seems therefore that both claims are unwarranted. Superordinate structure is not all-important in either projecting a property amongst categories or in mediating centrality effects. These two claims about superordinate structure will be revisited in Experiments 10 and 11 of Chapter IV.

3.9.3 Vagueness and directionality

Experiment 5 investigated for centrality effects under vague conditions. Both the candidate properties (substance Z, and substance W) as well as the premise and conclusion categories (Mammal 1, or Mammal 2) were left vague. The results showed that centrality operates under vagueness. Structural-mapping models cannot capture this finding since for such models to work predicates, as well as their relations to other predicates, need all be clearly specified.

Yet, one could claim that the participants translated the vague properties as concrete ones, and reasoned from those. Even if we grant that participants reasoned about concrete properties, the structural-mapping model still fails to predict a preference in projecting the central over the less-central properties (the depended-on over the dependent properties). That is because it does not consider the status of a property in a relation (refer to the General Discussion of Chapter II). However, a preference to generalize central properties was evidenced in all our studies. In Experiments 3 and 4 participants preferred to project causes over effects in the causal condition, rather than their counterparts in the control condition. Further, in experiments 6 and 7 participants gave overall higher inductive strength estimates for the depended-on rather than for the dependent properties.

3.9.4 Feature centrality: Centrality in a causal or a dependency space?

Ahn and Lassaline (1995) claim that feature centrality is determined by virtue of a feature's causal relations (see also Ahn and Dennis, 1997). On this account, the centrality of a feature is a function of a feature's causal status. Our hypothesis extends this notion to asymmetric dependency relations of which causal ones are but a special case. In Experiment 5

the relative centrality of two properties was defined by embedding them in an unspecified dependency relation. The results showed that participants preferred to project the central over the less-central properties much in the same way as with causal relations (see Experiments 3 and 4). This supports the view that feature centrality is not confined to causal status, but more generally to the status of a property in a concept's dependency space. (For independent evidence see Sloman et al. 1998, Studies 3A and 3B)

3.9.5 Blankety blank

Were the candidate properties used in the present experiments blank? Certainly the properties of Experiments 6 and 7 (and the abstract properties of Experiment 5), such as neurotransmitter octopamine, have the same flavor of unfamiliarity as the properties that Osherson et al. (1990) and Sloman (1993) used. At the same time, our studies provided information about how the candidate properties were related to each other, or how each property was related to a premise category's functions. Thus the properties were unfamiliar in a sense (most people are unfamiliar with the neurotransmitter octopamine), but not in terms of relative centrality status. Whether such predicates qualify as blank, depends on where a theorist stands on the centrality issue. For theorists who assume that centrality information (information about the dependency status of a property) does not influence inductive reasoning, then such properties are blank. Our whole theory centers on the belief that centrality information does influence projectibility, and hence we do not qualify such predicates as blank.

To sum up so far, experiments in Chapter II showed that the more immutable a property for an animal category the higher its generalizability to other animal categories. The results also suggested that immutability and centrality in a dependency space measure the same underlying factor (presumably conceptual centrality). Experiments in Chapter III built on this finding. They searched for centrality effects with centrality defined from a single dependency chain. The results showed that centrality defined as such influences projectibility

and that shared structure does not interact with centrality. A centrality effect was also shown under conditions of vagueness, i.e. with unspecified dependency relations. This implies that (1) current structural mapping models will prove unable to account for centrality effects, and (2) that feature centrality is best conceived as centrality in a dependency space (as we suppose) rather than as centrality in a causal space (as the causal status hypothesis suggests). Taken together, the findings support the Sloman et al. (1998) notion that centrality can be ultimately defined in terms of local pairwise dependency relations. Chapter IV continues to examine the centrality hypothesis by focusing on another critical aspect of the Sloman et al. notion, i.e. that a feature is central for a concept to the extent that other features depend on it.

CHAPTER IV:
CENTRALITY AND PROJECTIBILITY FROM
THE "FEW/LOTS" DEFINITION

4.1 INTRODUCTION

Experiments of Chapter IV operationalize relative feature centrality from the "Few/Lots" definition and examine its influence on projectibility. We stipulate that lots of an animal's functions depend on a property but few on another property of the same type (e.g., lots of an elephant's properties depend on the enzyme elastase, but few on the enzyme aliesterase). Participants are asked either to rate the likelihood that a target species has each of these properties or else to choose the property that a target species is more likely to have. The property upon which lots of an animal's functions depend (e.g. elastase) is taken to be the central property, while the property upon which few of the animal functions depend (e.g. aliesterase) to be the less-central property. The following sections explain why the "Few/Lots" operationalization was used.

4.1.1 The "Few/Lots" definition captures Sloman et al.'s (1998) notion of feature centrality

The "Few/Lots" operationalization captures the spirit behind the Sloman et al. (1998) notion that a feature is central to the extent that other features depend on it. Although in the Sloman et al. model feature centrality ultimately derives from pairwise dependency relations, defining relative feature centrality from a single dependency chain (as in the experiments of Chapter III) is an extreme rather than a typical case. The "Few/Lots" definition also leads to the construction of more ecologically valid examples in the sense that most biological properties are intertwined with multiple other biological properties and functions. Evidence that centrality from the "Few/Lots" definition influences projectibility would further suggest that centrality in a concept's dependency space constrains property induction. This definition also helps to rule out an interpretation that the results of Chapter III afford, as shown in the next section.

4.1.2 The "Few/Lots" definition helps control for necessity and sufficiency

In Chapter III relative feature centrality was defined in terms of a single dependency chain, the depended-on properties being the central and the dependent the less-central properties, and its effect on property induction was examined. This operationalization helped contrast our predictions against those of structural-mapping models. Recall that most of the evidence for the principle of systematicity as an inductive constraint came from studies that compared the projectibility of features embedded in causal (or temporal) relations to ones that were not (see Lasaline, 1996; Wu & Gentner, 1998 [for a brief exposition of their studies refer to the introductory chapter]). The findings of Chapter III supported our claim. They showed that depended-on properties were more projectible than their dependent counterparts; central properties were more projectible than less-central ones.

One might attribute the results to a default assumption that depended-on properties are necessary but insufficient for the presence of the dependent properties. In the case of a cause and effect relation, for instance, the presence of a cause may be assumed as necessary but not sufficient for the presence of the effect. On this assumption, whenever the effect is present the cause is present as well (the cause is necessary), but there may be cases where the cause is present but the effect is absent (the cause is not sufficient). Under such an assumption, therefore, it normatively follows that "central" properties are more projectible than their "less-central" counterparts. Although it is not clear why there is or there should be such a default assumption, the "Few/Lots" definition investigates to what extent centrality effects can be explained by an appeal to necessity and sufficiency. What is manipulated in Chapter IV is the number of properties that depend upon a candidate feature, not the status of the candidate feature: both the central and the less-central features are depended-on. This manipulation should thus control for any default assumptions about necessity and sufficiency.

Note that we do not argue that beliefs about necessity and sufficiency leave property induction unaffected. Quite the contrary, all else being equal, stipulating that various features can account for the functions that depend upon a candidate feature, should reduce the

projectibility of that particular feature. That is because any one of those features would be sufficient to promote coherence. Stipulating, for example, that lots of elephants' enzymes help increase its metabolic rate reduces the belief that any particular enzyme is present in rhinos. Instead, the "Few/Lots" definition aims to rule out the possibility that the centrality effects can be fully explained by an appeal to default assumptions about necessity and sufficiency.

4.1.3 The "Few/Lots" definition offers a different manipulation of shared structure

So far, our experiments have failed to detect an interaction between (overall) shared structure and feature centrality on category-based inference. However, based on conceptual arguments we have claimed that the effect of centrality on projectibility should be mediated by the extent to which the base and target concepts share structure. That is because centrality is concept relative; e.g., *roundness* is a central property for the concept basketball but not for the concept cantaloupe (Medin & Shoben, 1988). In cases where the premise and conclusion categories share very different structure, the centrality effect may even be reversed. It is less likely, for instance, that humans and chairs share conceptually central rather than conceptually peripheral properties.

Why then our experiments have failed so far to detect such an interaction? A reason might be that shared structure was not manipulated widely enough; all premise and conclusion categories involved animals and therefore shared lots of structure. Another (and possibly complementary) reason might be because in many experiments (especially in Experiments 6 and 7) the dependent properties were concrete and fairly general such as metabolism (but notice that Experiment 5 used abstract properties). Thus, it could be safely assumed that all target animals shared these properties. The point is that what might mediate the effect of feature centrality on projectibility may not be the extent to which the base and target categories share overall structure (properties and relations), but rather the extent to which the target shares the specific properties that depend on the candidate feature.

The "Few/Lots" definition goes some way in addressing this issue by leaving the properties that depend upon the central and less-central candidate inferences vague. Intuitively, the more similar the base and target categories are, the more likely that the target will share the dependent properties specified for the base. All else equal, the more properties the target shares with the base, the greater the need to support them and hence the greater the need to project the central candidate inference. Following this line of reasoning, as the amount of shared structure decreases the preference to project the central property should decrease as well. Evidence therefore for a centrality by shared structure effect in experiments of the current chapter, would support that the effect of feature centrality on projectibility is mediated by whether the target is believed to share the base's properties that depend on the candidate feature.

Experiments 10 and 11 of Chapter IV, like Experiments 6 and 7 of Chapter III, manipulate superordinate structure. They use animal pairs high in similarity from the same superordinate, low in similarity from the same superordinate, and low in similarity from different superordinates. Hence, these experiments also address the role of superordinate structure in property induction and in mediating the effect of centrality.

4.1.4 The "Few/Lots" definition provides a test for vagueness

As discussed earlier, a precondition for structural mapping models to make any predictions is that predicates, as well as their relations to other predicates, be all completely specified. Stating that lots (or few) of an animal's functions depend on a particular feature without specifying which, leaves this condition unsatisfied. Evidence for asymmetries in the preference to project the central over the less-central properties would therefore fall outside the domain of such models.

4.2 EXPERIMENT 8: FAMILIAR CATEGORIES, ABSTRACT PROPERTIES

The purpose of Experiment 8 was twofold: (1) to provide evidence that centrality from the "Few/Lots" definition influences projectibility, and (2) to detect empirically a possible boundary condition for the centrality effect. Based on the findings of Chapter III, no interaction between feature centrality and shared structure was expected. However, if such an interaction is observed, it will support that the centrality effect gets mediated by the extent to which the target category shares (or more accurately, it is believed to share) the properties specified for the candidate feature for the premise.

METHOD

Participants. Twenty-four University of Durham students voluntarily participated in this experiment.

Design and Materials. Participants were informed that a mammal had two properties (Z and W) and that many of that mammal's physiological functions depended on Z (the central property), but only few on W (the less-central property). Participants were asked to choose the property (Z or W) that they believed another mammal was more likely to have. Sixteen items were constructed such that half included similar and half dissimilar pairs of mammals. Table 4.1 overleaf presents sample items from each shared structure condition.

Animal pairs of different similarity were selected from the Osherson et. al. (1990, p.196) study where they provide a pair-wise similarity matrix. However, for some items we relied on our intuitions. To check whether our similarity manipulation was justified, we gathered similarity estimates for the categories used. An independent group of twelve students from the University of Durham voluntarily rated the biological similarity of the 16 mammal pairs in a scale from 0 to 10, where 0 was labeled as "highly dissimilar" and 10 as "highly similar." For the High-Similarity condition the mean rating was 7.39 (SE = .43; min = 5.92; max = 8.92), whereas for the Low-Similarity condition the mean was 1.74 (SE = .18; min =

.92; max = 2.33). The results justify therefore the assignment of category pairs into High-Similarity and Low-Similarity conditions. The similarity data are presented in the accompanying floppy disk.

Table 4. 1 Sample items from each Shared Structure condition for Experiment 8.

Shared structure

High Similarity

Fact: Many of a horse's physiological functions depend on Z, but only a few depend on W.

Choose the statement below that you consider more likely to be true.

- A. Cows have Z. Central
B. Cows have W. Less central

Low Similarity

Fact: Many of a seal's physiological functions depend on Z, but only a few depend on W.

Choose the statement below that you consider more likely to be true.

- A. Horses have Z. Central
B. Horses have W. Less central
-

Two lists of materials were constructed, each consisting of eight high similarity items and eight low-similarity items. The two lists differed in that the order of the premise and conclusion categories in List 1 was reversed in List 2. For example, List 1 contained the premise-conclusion pair "horse-cow" whereas List 2 contained "cow-horse". Table 4.2 presents the full set of items of List 1. Twelve participants were assigned in each list. The presentation order of the items within each list was randomized for each participant. The order of the two response alternatives was pseudo-randomized for each participant; in each shared structure condition half the items presented the central response first, half the central response second.

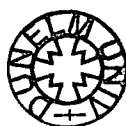
Table 4. 2 Animal pairs for Experiment 8, List 1.

High Similarity		Low Similarity	
Premise	Conclusion	Premise	Conclusion
Horse	Cow	Seal	Horse
Mouse	Squirrel	Cow	Mouse
Dolphin	Seal	Elephant	Beaver
Rhino	Elephant	Squirrel	Rhino
Chimpanzee	Gorilla	Gorilla	Deer
Lion	Tiger	Tiger	Chimpanzee
Deer	Antelope	Antelope	Dolphin
Beaver	Raccoon	Raccoon	Lion

Procedure. Participants were presented with a booklet, containing one of the resulting two lists of materials. There was no time restriction but participants were encouraged to work quickly. Participants received the following instructions:

Imagine that you are a first-year undergraduate in biology interested in the physiology of mammals. You recently got hold of an authoritative book on the physiology of mammals only to find out that some of the pages were missing or were torn apart. You are left with some excerpts from the book stating facts about the physiology of some mammals. You are informed that a given mammal has properties Z and W. Those properties may refer to any physiologically meaningful aspect of that mammal. You know nothing about the properties Z and W except that lots of the physiological functions of the given mammal depend on Z, but few depend on W. Beneath each excerpt you are presented with a different mammal and you are asked to choose the property (Z or W) that this mammal is more likely to have.

Before the experiment commenced participants worked through an example of the same format as the test items. They were also told that there was no right or wrong answer - they just had to choose the statement that they considered more likely to be true. Finally, they were instructed that the properties Z and W were not meant to refer to the same properties across items and that each item should be treated separately.



RESULTS

Coding. The choice of the central property (Z) was coded as 1 and the choice of the non-central property (W) as 0. For each participant, the number of central choices divided by the total number of choices was calculated for each similarity condition. Scores higher than .50 represent a positive centrality effect, scores lower than .50 a negative centrality effect, and scores close to .50 a null centrality effect.

The results are presented in Table 4.3. Overall, participants chose to project the central over the less-central property among high-similar but not among the low-similar pairs of mammals. That is, a positive centrality effect is shown for high similarity but not for low similarity items.

Table 4.3 Mean (SE) percentage choice of the central property.

Shared Structure	
High-Similarity	Low-Similarity
.95 (.02)	.44 (.09)

A 2 (List) x 2 (Shared structure) mixed ANOVA, with repeated measures on the last factor, showed a significant main effect of Shared structure ($F_1(1, 22) = 30.45, p < .001$; $F_2(1, 14) = 275.53, p < .001$). Both List and the List by Shared structure failed to reach significance (all $F_s < 1$). The mean estimate for the high-similarity condition was significantly greater than chance (participants: $t(23) = 19.94, p < .001$; items: $t(7) = 18.81, p < .001$; chance level = .50). The mean estimate for the low-similarity condition was not significantly different than chance for participants ($|t(23)| < 1$; chance level = .50), but it was for items ($t(7) = 3.24, p < .05$; chance level = .50).

DISCUSSION

Experiment 8 addressed whether feature centrality defined in terms of the "Few/Lots" definition influences property induction. In the case it did, it further aimed to detect a possible

boundary condition for this effect by manipulating the similarity between the premise and conclusion categories. The results showed a main effect of shared structure (on the effect of centrality) - high similarity items showed a positive centrality effect (mean score higher than .50), whereas low-similarity items showed no robust centrality effect (mean score close to .50). Otherwise stated, for the first time, the results showed an interaction between shared structure and centrality on projectibility - feature centrality influenced projectibility for similar but not for dissimilar pairs of mammals.

The findings of Experiment 8 should be interpreted with caution. One could object that the results are an artifact of a demand characteristic created by the forced choice procedure. If participants were given a wider choice of estimates (e.g., a “no preference” option) this result might disappear. (Notice, however, that such a claim is not supported from our studies; see e.g. Experiments 2 and 7) To remove this plausible interpretation, all our studies henceforth asked for likelihood estimates while putting the central and less-central properties side by side for evaluation. This meant that if participants were indifferent in projecting the central over the less-central properties, they were able to say so.

4.3 EXPERIMENT 9: UNSPECIFIED CATEGORIES, ABSTRACT PROPERTIES

Following up from Experiment 8, Experiment 9 aimed to test the reliability and generality of centrality effects by using the more sensitive measure of conditional likelihood estimates. Experiment 9 examined the effects of feature centrality when all that is known is the superordinate category, mammal or bird. Participants were asked to provide estimates for two items involving abstract properties. One item asked participants to project a central and a less-central property between mammals, while the other between a bird and a mammal.

METHOD

Participants. Twenty-four University of Durham students voluntarily participated in this experiment.

Design and Materials. Experiment 9 crossed Centrality (central or less-central properties) with Shared structure (same-superordinate or different-superordinate animal pairs) in a repeated measures design on both factors. Because the experiment involved only two items, no F_2 analyses were carried out.

The participants provided likelihood estimates for the two items below (the bold-faced headings are for clarification purposes and did not appear in the actual questionnaire):

A. The Mammal 1 → Mammal 2 Item (M → M)

Fact: Many of mammal's 1 physiological functions depend on property X, but very few depend on property Y.

Rate the likelihood of the following statements.

A. Mammal 2 has X. _____%

B. Mammal 2 has Y. _____%

B. The Bird 1 → Mammal 2 Item (B → M)

Fact: Many of bird's 1 physiological functions depend on property Z, but very few depend on property W.

Rate the likelihood of the following statements.

A. Mammal 2 has Z. _____%

B. Mammal 2 has W. _____%

Procedure. Participants were asked to provide likelihood estimates for the two items above which appeared in a single page. They were not explicitly told that mammal 1 and mammal 2 were meant to be different, although it is assumed that participants treated them as different. The order of presentation of the two items in the questionnaire, as well as the order of the two response alternatives within an item was counterbalanced. Properties X and Z always referred to the central properties whereas properties Y and W to the less-central ones. Because all properties were abstract we did not consider it necessary to counterbalance for the names of the central and the less-central properties.

RESULTS AND DISCUSSION

The results are summarized in Table 4.4. A centrality effect is detected for the M→M item but not for the B→M item. Further, similarity seems to strongly influence the projectibility of the central but not the less-central properties.

Table 4. 4 Mean (SE) inductive strength estimates for Experiment 9.

Centrality	Shared structure	
	M→M Item	B→M Item
<u>Central Property</u>	.73 (.04)	.47 (.05)
<u>Less-central Property</u>	.53 (.06)	.47 (.05)

A 2 (Shared structure) x 2 (Centrality) repeated measures ANOVA showed significant main effects of both Shared structure ($F_1(1,23)=9.76, p<.01$) and Centrality ($F_1(1, 23)=7.82, p<.01$). The Shared structure x Centrality interaction was also significant ($F_1(1, 23)=5.28, p<.05$). For the M→M item moving from the less-central to the central property led to a significant increase in inductive strength of .20 ($t(23)=3.35, p<.005$). For the B→M item level of centrality had a null effect. For the central property moving from the M→M to the B→M item led to a significant drop in inductive strength of .26 ($t(23) = 4.51, p<.001$), whereas for the less-central property the same movement resulted in a non-significant decrease of only .06 ($t(23)<1$).

In sum, Experiment 9, using the more sensitive measure of likelihood estimates, showed that central properties are indeed more projectible than less-central ones but only across categories that share a common superordinate. The results rule out therefore the interpretation that the centrality effect of Experiment 8 was an artifact of the forced-choice procedure. If participants were indifferent in projecting either property they could assign to both properties the same inductive strength estimate. Furthermore, similarity strongly influenced the participants' willingness to generalize the central property but only weakly their willingness to generalize the less-central property.

Taken together experiments 8 and 9 suggest a strong effect of feature centrality for same-superordinate highly similar animals (similar mammals) and no effects for different-superordinate dissimilar animals (Experiment 9). Experiment 9 did not directly assess the issue of whether same-superordinate dissimilar animals can give rise to centrality effects, though Experiment 8 suggests that they do not. Following up from these experiments, Experiment 10 further examines the boundary conditions of the centrality effect by using animal categories that are highly similar from the same superordinate, dissimilar from the same superordinate, and dissimilar from a different superordinate.

4.4 EXPERIMENT 10: FAMILIAR CATEGORIES, CONCRETE PROPERTIES

Experiment 10 examined centrality effects in a more ecologically valid context where both the animal categories are specified and the properties are meaningful (though unfamiliar). A boundary condition for the centrality effect on projectibility was investigated by crossing Centrality (central vs. less-central properties) with Shared structure (high-similarity same-superordinate vs. low-similarity same-superordinate vs. low-similarity different-superordinate animal pairs) in a repeated measures design on both factors. Based on the evidence from Experiment 9, we expected no centrality effects in the different-superordinate low-similarity condition (LSDS). Experiment 8 suggests no effect of feature centrality in the same-superordinate low-similarity condition (LSSS), though it employed the less powerful forced-choice procedure.

METHOD

Participants. Twenty-four University of Durham students volunteered to participate in the experiment.

Design and Materials. Experiment 10 crossed Centrality (central or less central properties) with Shared structure (high-similar same-superordinate (HSSS), or low-similar same

superordinate (LSSS), or low-similar different-superordinate (LSDS) animal pairs) in a repeated measures design on both factors.

Participants were informed that an animal had two properties: one upon which lots of its functions depended (the central property), and one upon which only few of its functions depended (the less-central property). Participants were asked to estimate the likelihood that another animal had each of these properties.

The same animal pairs and properties were used as in Experiment 6, but the wording of the items was different. Eighteen items were constructed, 6 in each shared structure condition. Table 4.5 presents sample items from each shared structure by centrality condition.

Table 4.5 Sample items from each Shared Structure by Centrality condition of Experiment 10.

Shared structure

HSSS

Fact: Many of a squirrel's physiological functions depend on the enzyme amylase, but only a few depend on the enzyme streptokinase.

Please rate the likelihood of the following statements.

- A. Mice have amylase. _____% Central
B. Mice have streptokinase. _____% Less central

LSSS

Fact: Many of a bear's physiological functions depend on the enzyme amylase, but only a few depend on the enzyme streptokinase.

Please rate the likelihood of the following statements.

- A. Mice have amylase. _____% Central
B. Mice have streptokinase. _____% Less central

LSDS

Fact: Many of an eagle's physiological functions depend on the enzyme amylase, but only a few depend on the enzyme streptokinase.

Please rate the likelihood of the following statements.

- A. Mice have amylase. _____% Central
B. Mice have streptokinase. _____% Less central
-

Three types of properties were used: enzymes, hormones, and neurotransmitters. For any given item, both properties were of the same type (e.g. both were enzymes as in the illustration above). Six lists of materials were constructed by counterbalancing the assignment of property type to animal pairs. Table 4.6 shows the design of the 6 lists of materials.

Table 4. 6 Design of the 6 lists of materials of Experiment 10.

Triple #	List1	List2	List3	List4	List5	List6
1	EHN	EHN	HNE	HNE	NEH	NEH
2	EHN	EHN	HNE	HNE	NEH	NEH
3	HNE	NEH	EHN	NEH	EHN	HNE
4	HNE	NEH	EHN	NEH	EHN	HNE
5	NEH	HNE	NEH	EHN	HNE	EHN
6	NEH	HNE	NEH	EHN	HNE	EHN

Note. EHN = HSSS enzyme, LSSS hormone, LSDS neurotransmitter.
HNE = HSSS hormone, LSSS neurotransmitter, LSDS enzyme.
NEH = HSSS neurotransmitter, LSSS enzyme, LSDS hormone.

Each of the six lists had a counterpart where the names of the central and the less central properties were reversed, making 12 lists of materials in all.

Procedure. Participants were presented with a booklet, containing one of the resulting 12 lists of materials. There was no time limit but participants were encouraged to work quickly. Experiment 10 used similar instructions to the ones used in Experiment 8, with minimal changes to compensate for the use of unfamiliar properties and animals from different superordinates. Specifically, “animal” was substituted for “mammal,” “two properties” for “properties Z and W,” and “(Z or W)” and the sentence “You know nothing...on W” were both deleted from the instructions. Prior to the experiment, participants worked through an example of similar format to the test items.

RESULTS

Table 4.7 summarizes the results. The centrality effect seems proportional to shared structure. This seems to be because as one moves from the HSSS to the LSSS to the LSDS condition, the inductive strength estimates for the central properties decrease steeply, whereas the estimates for the less-central properties decrease to a lesser extent.

Table 4.7 Mean (SE) inductive strength estimates as a function of Centrality and Shared structure for Experiment 10.

Shared structure	Centrality	
	Central property	Less-Central property
<u>HSSS</u>	0.75 (.03)	0.55 (.05)
<u>LSSS</u>	0.53 (.04)	0.47 (.04)
<u>LSDS</u>	0.44 (.04)	0.47 (.05)

A preliminary test showed no effects of the type of list used. The data were therefore collapsed across lists. Two 3 (Shared structure) x 2 (Centrality) analyses of variance were carried out. In the F_1 analysis both factors were repeated measures, in the F_2 analysis only Centrality was repeated measures. The analyses showed a main effect of Shared structure ($F_1^*(2, 22) = 9.21, p < .005$; $F_2(2, 15) = 50.84, p < .001$) and a main effect of Centrality ($F_1^*(1, 23) = 11.51, p < .005$; $F_2(1, 15) = 24.23, p < .001$). The Shared structure x Centrality interaction was also significant ($F_1^*(2, 22) = 5.65, p < .05$; $F_2(2, 15) = 20.75, p < .001$).

To examine the locus of the centrality effect, multiple pairwise t-tests were carried out across participants. For the HSSS condition, moving from the less-central to the central properties was accompanied by a significant increase in inductive strength of .20; ($t(23) = 4.46, p < .005$). For the LSSS condition, a marginally significant centrality effect of .08 was detected ($t(23) = 1.92, p < .07$). For the LSDS condition a slightly negative but non-significant effect was detected ($|t| < 1$). The results show therefore a strong effect of centrality in the HSSS condition, a marginal effect in the LSSS condition and no effect in the LSDS condition.

The marginal effect of Shared structure in the LSSS condition leaves open the question of whether a common superordinate is sufficient for centrality effects. We will return to this claim below.

The Shared structure by Centrality interaction seems to be caused by the fact that shared structure differentially affected the generalizability of the central and the less central properties. The projectibility of the central properties was influenced by shared structure condition ($F_1^*(2, 22) = 22.67, p < .001$; $F_2(2, 15) = 65.16, p < .001$), whereas the projectibility of the less-central properties was not significantly influenced by shared structure condition for participants ($F_1^*(2, 22) = 1.59, p > .23$), but it was for items ($F_2(2, 15) = 5.47, p < .05$).

Our experiments have yet to provide a fine-grained quantification of the relation between similarity and inductive strength estimates for the central and the less-central properties. Taken together, Experiments 8, 9, and 10 show little change in the mean inductive strength for the less-central properties as one moves across similarity conditions - they are all close to .50 (see Tables 4.2, 4.3, and 4.6). In search for a fine grained quantification of this relation, the correlation between the mean similarity ratings and the corresponding mean inductive strength estimates for each item were computed. The correlation between similarity and inductive strength for central properties was .98, and between similarity and inductive strength for less-central properties was .61. Their difference was found significant by a t-test for non-independent *r*s developed by Williams in 1959 ($t(15) = 5.02, p < .01$). The correlation between Similarity and the difference in inductive strength between the central and the less central items was .90, $p < .001$. In sum, the results show that similarity is (1) an excellent predictor of the projectibility of central properties, (2) a somewhat worse predictor of the projectibility of the less central, and (3) a good predictor of the centrality effect. Table 4.8 presents the correlation statistics for the present Experiment as well as for Experiments 6, 7, and 11.

Table 4. 8 Correlation statistics for items between estimates of shared structure and estimates of inductive strength for central properties, less central properties, and their difference.

	Experiment #			
	10	11	6	7
Correlation between estimates of shared structure and estimates for:				
Central properties	.98**	.98**	.91**	.94**
Less-central properties	.61*	.93**	.93**	.97**
Central - Less-central properties	.90**	.94**	-.13	-.02

Note. * $p < .05$; ** $p < .01$.

Role of superordinate structure. To examine the claim that sharing a superordinate decides inductive strength, the data were collapsed across centrality conditions and unplanned multiple pairwise comparisons were carried out using Tukey's HSD test. Across both participants and items, the mean for the HSSS condition was significantly higher than that for either the LSSS or the LSDS conditions. The means for the LSSS and LSDS conditions were not significantly different from each other. (For participants: mean HSSS=64.82; mean LSSS=50.09; mean LSDS=45.66; $CD=8.57$ [$MS_{error}=148.81$; $n=24$; $q_r(3, 46) = 3.44$]).

To examine the claim that sharing a superordinate mediates the effect of centrality on induction, six Shared structure by Centrality analyses of variance were carried out (three for participants and three for items) and the interactions were observed. Each of these analyses of variance contrasted two shared structure conditions. A significant Shared structure by Centrality effect was detected for the analyses contrasting the HSSS and the LSSS conditions ($F_1(1,23)=10.75$, $p < .001$; $F_2(1,10)=11.30$, $p < .01$), and the HSSS and the LSDS conditions ($F_1(1,23)=10.20$, $p < .005$; $F_2(1,10)=40.36$, $p < .001$). For the analyses contrasting the LSSS and LSDS conditions the interaction was significant for items ($F_2(1,10)=9.70$, $p < .05$) and marginally significant for participants ($F_1(1,23)=3.71$, $p < .07$).

DISCUSSION

In Experiment 10 the effect of centrality on property induction was found proportional to the amount of shared structure between the premise and conclusion categories. A significant

centrality effect was reported for same-superordinate high-similar animal pairs, a marginal effect for same-superordinate low-similar pairs, and no effect for different-superordinate low-similar pairs. This finding was due to that shared structure influenced more the projectibility of the central properties than it influenced the projectibility of the less central properties. Correlation analyses across items between similarity estimates and inductive strength estimates showed a .98 correlation for central properties, a .61 correlation for less-central properties, and a .90 correlation for the difference between the inductive strength of the two centrality conditions.

The findings of Experiment 10, as those of Experiments 6 and 7, suggest that superordinate structure is not all-important in estimates of inductive strength or in mediating the effect of centrality. In the view of the empirical evidence, it is more accurate to state that both estimates are proportional to judged similarity; the higher the similarity between two categories the higher the argument strength, and the higher the difference between the projectibility of the central and the less central properties. Taken together experiments 8, 9, and 10 show that the effect of feature centrality, as defined from the "Few/Lots" definition, is proportional to similarity.

4.5 *EXPERIMENT 11: CENTRAL VERSUS LESS-CENTRAL VERSUS NO INFORMATION GIVEN PROPERTIES*

A pretty robust result across the experiments of Chapter IV is that the more central a property the higher its projectibility across categories in proportion to the degree to which they are similar. This finding directly supports our initial hypothesis (see Chapter I). A close inspection of the results, however, reveals that the locus of the shared structure by centrality interaction is that similarity influences the projectibility of the central but not of the less-central properties (notice though that the correlation analyses of Experiment 10 showed that the projectibility of the less-central properties significantly covaries with similarity though to a lesser extent than the projectibility of the central properties). A parsimonious interpretation

of the findings is that they are not due to differences in conceptual centrality (as we claim) but rather due to similarity being a good surrogate for estimating the likelihood that a target category has the "central" but not the "less-central" properties. This is because lots of functions of a base animal depend on the "central" property, while few on the "less-central" one. In contrast, the locus of our explanation is that central properties are more projectible than less-central ones across categories (to the extent that these categories share structure) because they promote more coherence to the target concept. [Note that our hypothesis is supported by the experiments of Chapter III.]

One way that Experiment 11 investigated this possibility is by including a no-information-given condition. We reasoned that a property about which no centrality information is given might be a better control than a property for which it is stipulated that only few of an animal's properties depend on. If, as we propose, people are more confident in projecting central properties because these promote more coherence to the target concept, then we should observe a higher willingness to project central rather than no-information-given properties. A second way that Experiment 11 examined this possibility, is by comparing the correlation between similarity estimates and inductive strength estimates for the central properties to that between similarity estimates and inductive strength estimates for the less central properties. To the extent that effect of shared structure on projectibility is comparable for central and less central properties, we expected no difference between the two correlation coefficients.

METHOD

Participants. Thirty-four University of Durham first-year undergraduates in psychology volunteered to participate.

Design and Materials. Shared-structure (HSSS vs. LSSS vs. LSDS) was crossed with Centrality (central vs. less-central vs. no-information-given properties) in a repeated measures design on both factors.

The first three triples from Experiment 6 were used. Each participant had to rate nine items, three in each shared structure condition. Each item asked for three likelihood estimates, one in each centrality condition. Participants therefore had to provide 27 likelihood estimates in all. Table 4.9 presents a sample item from each Shared structure by centrality condition constructed from the second premise triple of Table 3.6.

Table 4.9 Sample items from each Shared Structure by Centrality condition of Experiment 11.

Shared structure

HSSS

Fact: Squirrels have:
 The neurotransmitter taurine upon which lots of their physiological functions depend.
 The neurotransmitter glycine upon which few of their physiological functions depend.
 The neurotransmitter tyrosine about which you have no information.

Please rate the likelihood of the following statements:

Mice have taurine. _____%	<u>Central</u>
Mice have glycine. _____%	<u>Less-Central</u>
Mice have tyrosine. _____%	<u>No-info</u>

LSSS

Fact: Bears have:
 The enzyme aliesterase upon which lots of their physiological functions depend.
 The enzyme streptokinase upon which few of their physiological functions depend.
 The enzyme elastase about which you have no information.

Please rate the likelihood of the following statements:

Mice have aliesterase. _____%	<u>Central</u>
Mice have streptokinase. _____%	<u>Less-Central</u>
Mice have elastase. _____%	<u>No-info</u>

LSDS

Fact: Eagles have:
 The hormone MSH upon which lots of their physiological functions depend.
 The hormone ACTH upon which few of their physiological functions depend.
 The hormone TSH about which you have no information.

Please rate the likelihood of the following statements:

Mice have MSH. _____%	<u>Central</u>
Mice have ACTH. _____%	<u>Less-Central</u>
Mice have TSH. _____%	<u>No-info</u>

As in Experiments 6 and 10, three types of properties were used: enzymes, hormones, and neurotransmitters. For a given item all properties were of the same type (e.g. see Table 4.8). Three lists of materials were constructed that counterbalanced the animal pair property-

type assignment in a Latin square design. Each of these lists had a counterpart where the names of the central and the less-central properties were reversed, making 6 lists in all.

Procedure. Participants were given a booklet containing one of the six lists of materials. The instructions of Experiment 10 were minimally revised to compensate for the changes in the present experiment; “two properties” was substituted by “three properties,” and participants worked through examples similar to the test items.

RESULTS

The results are summarized in Table 4.10. Central properties were the most projectible, followed by the less-central and the no-information-given properties. There seems to be a significant interaction between shared structure and central vs. either less-central or no-information given properties, but no interaction between shared structure and less-central vs. no-information given properties.

Table 4. 10 Mean (SE) inductive strength as a function of Similarity and Centrality for Experiment 11.

	Centrality		
	Central property	Less Central property	No-information property
Shared structure			
<u>HSSS</u>	80.43 (2.65)	62.00 (4.14)	57.74 (4.19)
<u>LSSS</u>	60.98 (4.50)	52.30 (3.85)	47.92 (3.90)
<u>LSDS</u>	55.76 (4.45)	51.64 (4.10)	47.83 (4.28)

Preliminary tests showed that there was no effect of list. The data were therefore collapsed across lists. A 3 (Shared structure) x 3 (Centrality) ANOVA was conducted across participants with repeated measures on both factors. F_2 statistics were not calculated because there were only three items per shared structure condition. A main effect of Shared structure was detected ($F_1^*(2, 32) = 12.88, p < .001$), and a main effect of Centrality ($F_1^*(2, 32) = 9.25,$

$p < .001$). The Shared structure x Centrality interaction was also significant ($F_1^*(4, 30) = 5.53$, $p < .005$).

Central versus less-central properties. A 3 (Shared structure) x 2 (central vs. less-central properties) evidenced a main effect of shared structure ($F_1^*(2, 32) = 15.87$, $p < .001$), a main effect of centrality condition ($F_1^*(1, 33) = 14.98$, $p < .001$), and a significant interaction ($F_1^*(2, 32) = 6.15$, $p < .01$). Central properties were more projectible than less-central ones in proportion to similarity.

The correlation coefficients between estimates of similarity and estimates of inductive strength for central properties and less central properties are presented in Table 4.8. Both correlation coefficients were high (.98 and .93 respectively), and not statistically different from one another.

To examine the claim that sharing a superordinate decides inductive strength, the data were collapsed across centrality conditions and unplanned multiple pairwise comparisons were carried out using Tukey's HSD test across participants. As in Experiment 10, the mean for the HSSS condition was significantly higher than that for either the LSSS or the LSDS conditions. The means for the LSSS and LSDS conditions were not significantly different from one another. (mean HSSS=66.72; mean LSSS=53.73; mean LSDS=51.74; CD=6.04 [$MS_{\text{error}}=107.24$; $n=34$; $q_r(3, 66) = 3.40$]).

To examine the claim that sharing a superordinate mediates the effect of centrality on induction, three Shared structure by Centrality analyses of variance were carried out for participants and the interactions were observed. Each of these analyses of variance contrasted two shared structure conditions. A significant Shared structure by Centrality effect was detected in all analyses (HSSS-LSSS: $F_1(1,33)=6.89$, $p < .05$; HSSS-LSDS: $F_1(1,33)=11.94$, $p < .005$; LSSS-LSDS: ($F_1(1,33)=5.27$, $p < .05$).

Central versus no-information-given properties. A 3 (Shared structure) x 2 (central vs. no-information-given properties) evidenced a main effect of Share structure ($F_1^*(2, 32) = 14.91$, $p < .001$), a main effect of centrality ($F_1^*(1, 33) = 18.97$, $p < .001$), and a significant

interaction ($F_1^*(2, 32)=11.22, p<.001$). Central properties were more projectible than no-information-given properties in proportion to similarity.

Less-central versus no-information-given properties. A 3 (Shared structure) x 2 (Less-central vs. no-information-given properties) evidenced a main effect of Shared structure ($F_1^*(2, 32)=5.53, p<.01$), and a main effect of centrality ($F_1^*(1, 33)=4.59, p<.05$). There was no Shared structure x Centrality interaction ($F_1^*(2, 32)<1$). Less-central properties were more projectible than no-information-given properties but this was unaffected by similarity condition.

DISCUSSION

The main objective of Experiment 11 was to detect whether the results of experiments 8, 9 and 10, that the more central a property the higher its projectibility among categories that share structure, were because central properties lend more coherence to the target concept, or else because the description of the central but not the less-central properties invited to use similarity as a surrogate of projectibility. To decide among these hypotheses we used similar materials as in Experiment 10 but we included a no-information-given condition. The rationale was that the no-information-given condition should provide a better control for our hypothesis. The centrality effect of Experiment 10 was replicated irrespective of which condition is chosen as a control. Central properties were more projectible than either less-central or else no-information-given properties in proportion to shared structure. This suggests that, as we assume, central properties are more projectible because they lend more coherence to the target concept.

Our account is further reinforced from the correlation analyses between similarity estimates and inductive strength estimates for the central and the less central properties. Both correlation coefficients were high (.98 and .93 respectively) and not significantly different from one another. In contrast to Experiment 10, Experiment 11 showed a comparable effect of shared structure on the projectibility of the central and the less central properties (see Table 4.8).

The findings of Experiment 11, as those of Experiments 10, suggest that superordinate structure is not all-important in estimates of inductive strength or in mediating the effect of centrality. There was a significant difference between the projectibility of properties among HSSS and either LSSS or LSDS categories. Properties were not projected differentially between LSSS and LSDS categories. Also, a significant Shared structure by Centrality interaction was shown from comparing any two shared structure conditions.

A rather counterintuitive result was that less-central properties were significantly more projectible than the no-information-given properties. Since no-information was given about these properties one might expect them to fall somewhere between the central and the less-central ones. Possibly putting all types of properties side-by-side for evaluation made the less-central property seem superior since more information was given for it.

4.6 GENERAL DISCUSSION

Experiments 8, 9, and 10 in Chapter IV defined relative feature centrality from the "Few/Lots" definition and showed that central properties are more projectible than less central ones in proportion to shared structure between the premise-conclusion categories. Experiment 11 ruled out the interpretation that this finding was because the inductive strength for the central but not for the less central properties covaried with estimates of similarity. First, the inductive strength for both types of properties covaried with similarity to the same extent. Second, a similar shared structure by centrality effect was shown when the control condition involved properties about which no information was given. The significant centrality effect shown in Chapter IV provides a non-trivial extension of the results of Chapter III, since the present definition better captured the spirit of conceptual centrality behind the Sloman et al. (1998) model; i.e., that central features are those upon which many other features depend.

As stated in the introduction, the "Few/Lots" definition also helped clarify whether the lack of interaction between shared structure and centrality in Chapter III was because (1)

similarity was not widely varied, or because (2) what mediates the effects of centrality on projectibility is the extent to which the target shares the properties that are specified for the base by the candidate feature. Since all experiments of Chapter IV showed such an interaction, claim (2) is supported. The effect of feature centrality on projectibility seems to be mediated therefore not by some measure of overall similarity but rather by shared structure with respect to the dependent properties. According to this interpretation, the experiments of Chapter III failed to show such an interaction because the dependent properties were all fairly general. It could thus be safely assumed that all the target categories had them. In the experiments of Chapter IV the dependent properties were left abstract, and hence it could be reasoned that the target category had them to the extent that it was similar to the base.

Experiments 10 and 11 by manipulating superordinate structure addressed two claims concerning the role of taxonomies in property induction: (1) that sharing a superordinate decides the strength of arguments, and (2) that sharing a superordinate mediates the centrality effect on property induction. If claim (1) were true, then we should not observe differences in property generalizability among HSSS and LSSS categories, but we should observe differences among HSSS or LSSS and LSDS categories. Contrary to claim (1), the results showed significant differences between the HSSS and the LSSS or LSDS conditions only mirroring the findings from Experiments 6 and 7 of Chapter III. If claim (2) were true, then we should not observe differences in the generalizability of the central and the less central properties among same superordinate categories, but we should observe differences among different superordinate categories. Contrary to claim (2), differences were detected (i.e. significant Shared structure by Centrality interactions) in all shared structure conditions. It seems therefore that both claims are unwarranted. Superordinate structure is not all important in either projecting a property amongst categories or in mediating centrality effects. The role of superordinate structure in property induction and in mediating the centrality effect will be discussed in the Final Discussion.

The "Few/Lots" definition also helped to rule out an alternative interpretation afforded for the results of Chapter III. Those experiments found that central features were

more projectible than less central ones. Because relative centrality was defined in terms of an asymmetric dependency chain, the results were open to interpretations in terms of necessity and sufficiency. If participants had interpreted the central (depended-on) properties as necessary but not sufficient for the presence of the less-central (dependent) properties, then the preference to project central properties can be described as rule following. In the experiments of Chapter IV both the central and the less-central properties were depended-on; what varied was the number of functions that depended on each. Therefore, the present centrality effects cannot be accounted for by default assumptions about necessity and sufficiency.

Finally, the functions that depended upon the central and the less-central properties were left vague. As stated in the introduction, vagueness prevents structural mapping models (e.g., Gentner, 1983) from making any predictions. Thus, the present findings cannot be accounted for by such models.

CHAPTER V:
CENTRALITY VERSUS VARIABILITY:
JUDGMENTS OF FREQUENCY AND
PROJECTIBILITY

5.1 CONCEPTUAL CENTRALITY VERSUS VARIABILITY

Conceptually, the degree of immutability of a feature for a concept is orthogonal to its degree of homogeneity. Conceptual centrality and homogeneity are aspects of two distinct views one can adopt about concepts: the inside view and the outside view, respectively (cf. Tversky & Kahneman, 1983). Conceptual centrality measures the extent to which a feature promotes coherence to a concept, the extent to which deleting it (or changing its value) affects the internal structure of the mental representation. In contrast, homogeneity measures the extent to which a property is stable across the remembered instances of a category. In sum, conceptual centrality refers to the degree to which a feature is mentally transformable while retaining coherence, whereas homogeneity is a statistical property that refers to the degree to which a property is transformable across actual instances of a category.

Immutability and homogeneity are dissociable. A feature can be mutable yet homogeneous due to happenstance. For ravens, for instance, being black is a homogeneous property (since almost all ravens are black) albeit one of their non-central properties (since imagining a raven that is not black does not require much mutation). Even better, *being white* is a very homogeneous property of refrigerators, but one of their less central properties. In the other end, a feature can be intrinsically variable yet immutable. A chameleon's ability to change color to disguise in its surroundings is one of its central properties (example seen in Keil, Smith, Simons, & Levin, 1998).

The examples above illustrate that, in theory, the inside and outside views may diverge. However, because of their common reference, the inside and outside views are usually compatible. Thus, measures of immutability and homogeneity converge (for empirical evidence see Sloman et al., 1998). For example, for most birds *having wings* is both an immutable and highly homogeneous property. Parenthetically, Keil et al. (1998) suggested that the high correlation between typicality and conceptual centrality may underlie the apparent success of the Roschean view that probabilistic representations are sufficient to explain all of categorization. In general, the difference between centrality and variability

closely maps into the difference between explanatory and statistical information. This point will be elaborated in the General Discussion at the end of this chapter.

Moving away from conceptual arguments, evidence suggests that variability and mutability are empirically dissociable. Medin and Shoben (1988) demonstrated that a feature judged equally typical for two concepts might be judged to have differential importance for deciding category membership. For example, although *roundness* was judged to be an equally typical feature of cantaloupes and basketballs, participants rated a square cantaloupe as a better new example of the category cantaloupes than a square basketball of the category basketballs. Similarity-to-an-ideal is one of the measures of immutability (Sloman et al., 1998). Hence, Medin and Shoben's study can be said to corroborate that homogeneity and immutability are empirically distinct.

Similarly to the Medin and Shoben (1988) study, our studies so far used candidate properties (albeit unfamiliar ones) that were equated in terms of homogeneity, all members of the base category were assumed to have the central and the less-central properties. Our studies demonstrated that central and less-central properties are differentially projectible. The results therefore show that centrality affects projectibility with the effect of variability controlled. Taking a step forward, the present studies pit variability against centrality information directly and look at judgments of projectibility and perceived variability. To the extent that conceptual centrality influences projectibility we expected the central and the less-central properties to be differentially generalizable, irrespective of the observed frequencies. Also, to the extent that centrality information directly influences perceived frequency, we expected the frequency of central features to be overestimated relatively to the frequency of the less-central properties.

To examine these hypotheses, participants were presented with 24 profiles of hippos. Profiles contained information about the features that a particular hippo had including two novel biological properties, a central and a less central one. The frequency of the central and the less central properties were manipulated such that the central property appeared less frequently than the less central property. Thus, one property was central and variable, the other less central and less variable. Participants were asked to recall how many out of the 24

hippos they saw had each of the properties (Frequency task), and to judge the likelihood that some target categories (e.g. rhinos) had each of these properties (Induction task). The order of these tasks was counterbalanced. To the extent that centrality exerts a greater influence on projectibility than variability, we expected higher inductive strength estimates for the central and variable property than for the less central and less variable property. Furthermore, to the extent that centrality information directly influences perceived frequency, we expected the frequency of central features to be overestimated relatively to the frequency of the less-central properties.

Both of our claims are motivated by the literature. Evidence that centrality influences projectibility comes from the studies reported so far. Suggestive evidence that centrality influences perceived frequency comes from Spalding and Ross (1994). Using artificial category learning stimuli they asked participants to rank features for their importance in category membership and also by their frequency. Features that were rated more important tended to be judged more frequent than those rated as less important, though they were not. Given that immutable features are more important (are weighted more heavily) than mutable features in categorical decisions (e.g., Ahn, 1999; Ahn and Lassaline, 1995; Ahn and Dennis, 1997; Sloman et al., 1998), this suggests that centrality influences perceived frequency.

Direct evidence that conceptual centrality influences frequency judgments comes from Sloman et al. (1998, Study 5). Participants were given a brief description of an indigenous people concerning their ceremonial procedure of becoming a hunter. To become a hunter an individual should first get a *tattoo* (central property) and then a *scar* (less-central property). Getting the tattoo was a necessary precondition for getting the scar. Participants saw 24 profiles of such individuals and they subsequently had to estimate how similar-to-an-ideal hunter was a hunter missing each of these features (Centrality task) and to judge the frequency of the properties in the profiles (Frequency task). The results showed that participants who completed the Centrality task first subsequently misjudged the frequency of the properties in the direction of centrality. So, the Sloman et al. study suggests that

conceptual centrality defined from a dependency chain influences judgments of frequency.

5.2 A SUBTLE MANIPULATION OF FEATURE CENTRALITY

In Experiment 12, the central property was a property upon which 5 properties depended (e.g. the hormone corticosterone helps the animal cope with stress). The less central property negated the properties stipulated for the central property (e.g. the hormone aldosterone does not help the animal cope with stress). This manipulation aimed to control for the relative salience of the central and the less central properties.

Experiments 13 and 14 used a subtler manipulation of centrality. Upon both the central and the less-central features the same number of properties and functions depended. However, the properties that depended upon the central feature were more central themselves. This relative definition of centrality derives from the Sloman et al. (1998) notion that a feature is central to the extent that other central features depend on it. We stipulated, for instance, that hippos have the substance neurine (the central property) upon which three properties about neurons depended, and the substance folline (the less-central property) upon which three properties about hair-follicles depended. All else equal, neurine is a more central property than folline because neurons are (presumably) more central than hair follicles.

This subtle centrality manipulation was aimed to clarify how centrality influences projectibility. In the experiments of Chapter III, where centrality was defined in terms of a single dependency chain and the dependent properties were all fairly general, centrality had a comparable effect in all similarity conditions. In the experiments of Chapter IV, where centrality was defined from the "Few/Lots" definition and the dependent properties were left vague, the effect of centrality was proportional to the similarity between the base and target categories. A joint interpretation of these results is that the effect of centrality is proportional to similarity to the extent that the target category is believed to share the dependent properties specified for the base. In the single dependency chain experiments, the depending on properties were all fairly general and hence no centrality by similarity interaction. In the

"Few/Lots" experiments, the depending on properties were left vague and hence the centrality by similarity interaction. Our hypothesis assumes, however, that the projectibility of a feature depends not only on feature overlap with respect to the dependent properties but also, critically, on the degree of centrality of the dependent properties. Experiments 13 and 14 directly address this prediction.

5.3 EXPERIMENT 12

Experiment 12 was a first attempt to assess the assumption that people reason from the "inside" - use conceptual information to guide their inferences - even when "outside" information is available. Participants were given two pieces of information about a property: (1) "inside information" about the centrality status of a property, and (2) "outside information" about the relative frequency of a property across instances of a category (variability). Based on these pieces of information, participants had to provide inductive strength and perceived frequency estimates. To the extent that participants reason from the "inside" we expected reliance on centrality to estimate inductive strength. Also, to the extent that centrality information influences perceived frequency, we should observe an overestimation of the judged frequency of the central property.

METHOD

Participants. Twenty-four students from the University of Durham participated voluntarily in this experiment.

Materials. Participants read 24 statements about hippos in which four different properties are described a variable number of times. After reading the 24 statements, participants completed a frequency task, in which they had to estimate the frequency of each of the four properties, and an induction task, in which they had to estimate the likelihood that each of four animals

(hippos, rhinos, mice, and falcons) possessed each of the four properties. The order of the tasks was counterbalanced, resulting in 2 lists of material.

Each statement about hippos appeared on a computer screen and contained information about 4 properties: 1) a central property, 2) a less-central property, 3) mud spots and 4) a heart. Underneath the central property five attributes were listed that depended on it. The attributes listed for the less central property negated those listed for the central property (see sample description below). The frequency of each property was manipulated such that the central hormone, the less-central hormone, the mud spots and the heart appeared on the screen in 15/24, 16/24, 17/24, and 24/24 descriptions respectively. Critically, the central property was slightly more variable than the less-central one. Whenever a property was not chosen by the computer program for a particular description, it did not appear in that description. For each participant, the program randomly generated 24 descriptions that fitted the frequency constraints mentioned above. Thus, each description appeared with a varying number of properties, from 1 to 4. The order of the properties within each description was randomized. A sample description that includes all properties (the central property in this description is corticosterone) is illustrated below.

This hippo has:

The hormone corticosterone that:

- 1) helps the animal cope with stress
- 2) regulates sleeping habits
- 3) reduces inflammation
- 4) increases blood flow to muscles
- 5) increases glucose metabolism

The hormone aldosterone that:

- 1) does not help the animal cope with stress
- 2) does not regulate sleeping habits
- 3) does not reduce inflammation
- 4) does not increase blood flow to muscles
- 5) does not increase glucose metabolism

Mud spots

A heart

As is seen in the sample description, the less-central condition was constructed by negating the properties listed for the central condition. This aimed to control for the relative salience of these properties as measured by the space allocated for each.

Design. The induction task data were analyzed with a 2 (Centrality) by 4 (Shared structure) ANOVA with repeated measures on both factors. The frequency task data were analyzed with a 2 (Centrality) by 2 (Task order) ANOVA with repeated measures on the first factor. No F_2 analyses were carried out since there was only one item per shared structure condition.

Procedure. After viewing the descriptions, participants were provided with a booklet containing one of the two lists of materials. One group of participants ($N=12$) received the Frequency task followed by the Induction task, whereas for another group of participants ($N=12$) the order was reversed. In each group, for half of the participants the central hormone was aldosterone and the less-central corticosterone, whereas for the other half the names were swapped. For the Frequency task, participants had to estimate the frequency of each property in the 24 descriptions they saw. For the Induction task, they had to rate the likelihood of an animal category (hippos, rhinos, mice, and falcons) possessing each of the four properties (Induction task). The four animals used in the induction task differed in their similarity from identical (hippos), very similar (rhinos), less similar (mice), and dissimilar (falcons). The order of the tasks was counterbalanced.

The Frequency task appeared in a single page of the booklet. The order of the four frequency questions was randomized for each participant separately. The Induction task appeared in four pages of the booklet, each page asking about the likelihood that a single animal category (e.g. rhinos) had each of the four properties. The order of the properties in each page was randomized for each participant separately. Also, the pages of the induction task were shuffled for each participant separately.

Participants were shown the following instructions:

You are about to see 24 partial descriptions of 24 different hippos. Each description contains some properties that the particular hippo has. Underneath some of these descriptions additional information is given concerning why these properties are useful (or not useful) for the animal. Please read the descriptions carefully and pay close attention to the properties of each hippo. Take as long as you like to study each description. When you feel you have familiarised yourself with the information, press the space bar to move to the next description. After you see all the descriptions, you will be asked some questions about them.

Frequency Task. In the Frequency task participants were asked to consider the 24 descriptions of the 24 different hippos they saw and to estimate how many out of these 24 hippos possessed each of the 4 properties (e.g., How many of the hippos you saw have mud spots? ____ / 24).

Induction Task. In the Induction task participants were asked to consider all the hippos (rhinos, mice, and falcons) in the world and to estimate, for each animal, the likelihood of possessing each of the 4 properties (e.g., “How likely do you think it is that rhinos have aldosterone? ____ %).

RESULTS

1. Central vs. Less-central: novel properties

Induction Task. The results are summarized in Table 5.1. Centrality had a strong effect in all shared structure conditions.

Table 5. 1 Mean (SE) inductive strength estimates for the novel properties of Experiment 12.

Shared structure	Centrality	
	Central property	Less-central property
<u>Hippos</u>	76 (3.46)	65 (5.02)
<u>Rhinos</u>	75 (3.44)	59 (5.22)
<u>Mice</u>	61 (5.27)	46 (4.29)
<u>Falcons</u>	66 (4.60)	50 (4.61)

Preliminary tests showed that neither the name of the hormones nor the order of the tasks influenced participants' inductive strength estimates. The data were therefore collapsed across these manipulations and F_1 analyses were carried out. No F_2 analyses were carried out because there was only one item per condition. A 4 (Shared structure) x 2 (Centrality) ANOVA with repeated measures on both factors revealed a main effect of Centrality ($F_1^*(1,23) = 12.81, p < .005$) and a main effect of Shared structure ($F_1^*(3,21) = 6.58, p < .005$). There was no significant Shared structure x Centrality interaction ($F_1^* < 1$).

Two-tailed pairwise t-tests showed significant centrality effects at all levels of Shared structure (for all, $p < .05$). On average, moving from the less-central to the central property resulted in a significant increase in inductive strength estimates of about 14.5 points.

As in the experiments of Chapter III, the failure to detect a Shared structure by Centrality interaction might be because the dependent properties were fairly general and thus applied equally well to all animals. That is, irrespective of level of shared structure, there was a comparable reason of preferring the central to the less-central properties. Strangely, the mean estimates were higher for falcons than for mice, although this difference was not significant.

Frequency Task. The results from the frequency task are presented in Table 5.2 below. Participants who received the frequency task second overestimated the actual frequency of the central property and underestimated the frequency of the less-central property. Participants who received the frequency task first judged the central and less-central properties as about equally frequent, though the less-central property was actually more frequent.

Preliminary tests showed that the name of the hormones did not affect the participants' judgments. The data were therefore collapsed across this manipulation. A 2 (Order of tasks) x 2 (Centrality) ANOVA was carried out for with repeated measures on the last factor. The analysis revealed a significant main effect of Centrality ($F(1, 22) = 8.77, p < .01$) and a significant Order x Centrality interaction ($F(1, 22) = 5.52, p < .05$). No main

effect of Order was found ($F < 1$). The mean (SE) estimated frequency for the central property was 16.96 (.57) and for the less-central property 15.75 (.71).

Table 5. 2 Actual frequency estimates and mean (SE) frequency estimates for each task order for Experiment 12.

	Central property	Less-central property
Actual Frequencies	15.00	16.00
Frequency task following induction.	17.33 (.79)	15.17 (1.08)
Frequency task preceding induction.	16.58 (.84)	16.33 (0.92)

The significant Order by Centrality interaction reveals a differential misjudgment of frequency depending on task order. Participants who completed the frequency task second significantly overestimated the frequency of the central property by about 2.30 points ($t(11) = 2.95$, $p < .05$, chance = 15). Participants who completed the frequency task first marginally overestimated the frequency of the central property ($t(11) = 1.89$, $p < .10$). We speculate that this is because participants who completed the induction task first focused on the dependency structure of the properties and used this as a guide to estimate the frequencies. Importantly, this result demonstrates that centrality can bias the perception of variability.

2. Familiar properties

Induction Task. Table 5.3 shows that the participants were willing to project mud spots to hippos and rhinos but not to falcons or mice. Participants' estimates might have been influenced either from direct experience or by relating mud spots to observable attributes that they are associated with, such as "living in lakes", and reason from those attributes. Heart was projected with maximum confidence. This validates that the participants understood the task.

Table 5. 3 Mean (SE) inductive strength estimates for the familiar properties of Experiment 12.

	<u>Mud Spots</u>	<u>Heart</u>
Shared structure		
<u>Hippos</u>	71 (4.13)	100 (.42)
<u>Rhinos</u>	52 (6.06)	100 (.04)
<u>Mice</u>	9 (3.53)	99 (.83)
<u>Falcons</u>	4 (2.31)	100 (.04)

Frequency Task. Table 5.4 shows that participants underestimated the frequency of both mud spots and heart by about 2.00 points. Possibly this is in part because the familiar properties were less salient than either the central or the less-central properties since they were allocated less space in the PC screen.

Table 5. 4 Actual frequency estimates and mean (SE) judged frequency estimates for the familiar properties of Experiment 12.

	<u>Mud Spots</u>	<u>Heart</u>
Actual frequency	17	24
Judged frequency	15.08 (.96)	21.88 (.70)

DISCUSSION

Experiment 12 examined the relative importance of centrality and variability in property induction with concrete but unfamiliar properties (hormones aldosterone and corticosterone). The results strongly suggest that participants were thinking in terms of dependency structure because: (i) centrality had a strong effect in inductive strength estimates, and (ii) participants overestimated the actual frequencies of the central properties when the frequency task came second. That is, both the inductive strength and frequency estimates were based on centrality rather than variability when both pieces of information were available. In our terms,

participants were thinking from the “inside.” Centrality had an effect in all levels of Shared structure which we attributed to the generality of the dependent properties.

The results are open, however, to a competing interpretation (we thank James Hampton for this suggestion). The less-central property was described in terms of negative relations, relations that are less memorable than positive ones. Hence, the frequency judgment data could be a result of people remembering more positive relations than negative ones. In turn, the perceived frequency judgments could have mediated the inductive inference data. In other words, the results could be attributed to thinking in terms of variability biased by the greater memorability of positive cues. Notice, however, that this interpretation cannot readily explain the observed Order by Centrality interaction in frequency estimates. Why did the frequency estimates for the central property increase when the frequency task came second?

5.4 EXPERIMENT 13

Experiment 13 aimed to rule out the possibility that the results of Experiment 12 can be fully explained in terms of memorability. Both the central and the less-central properties were described positively by stipulating that 3 properties depended upon each. However, the properties that depended upon the central feature were themselves more central than those that depended upon the less-central feature.

METHOD

Participants. A new sample of 24 University of Durham students volunteered to participate.

Design and Procedure. The design and procedure of the present experiment were very similar to those of Experiment 12. A difference was that in current experiment the central property was present in 14 out of 24 descriptions (in Experiment 12 it was present in 15 out of 24 descriptions). The frequencies for the less central property, mud spots, and heart remained unchanged; 16 out of 24, 17 out of 24, and 24 out of 24 descriptions, respectively. To make

the instructions clearer we added the sentence "Whenever a property does not appear in a description of a hippo, then that hippo does not have the property". Also the sentence "Underneath ... animal." was replaced by the sentence "Underneath some of these properties additional information is given about what functions depend on these properties."

Materials. The materials used were similar to the ones used in Experiment 12, except that the descriptions of hippos were different. A sample description that includes all properties (the central property in this description is aldosterone) is presented below:

This hippo has:

The hormone aldosterone that:

- 1) regulates the amount of oxygen in the blood
- 2) controls blood density
- 3) controls blood pressure

The hormone corticosterone that:

- 1) regulates the amount of green in the eye
- 2) controls eye colour reflectance
- 3) controls uniformity of eye colour

Mud spots

A heart

Notice that the type of dependency relations for the central and the less-central properties was also controlled for.

RESULTS

1. Central vs. Less-Central: novel properties

Induction task. The results are summarized in Table 5.5. The central property was more projectible than the less-central one. Shared structure seems to influence the projectibility of the less-central but not of the central property.

Table 5. 5 Mean inductive strength estimates for Experiment 13.

Shared structure	Centrality	
	Central property	Less-Central property
<u>Hippos</u>	80.58 (4.04)	74.13 (4.99)
<u>Rhinos</u>	83.13 (4.26)	67.04 (7.02)
<u>Mice</u>	80.58 (4.99)	52.29 (6.93)
<u>Falcons</u>	80.58 (5.10)	61.00 (5.63)

As in Experiment 12, preliminary tests showed that the name of the properties and the order of the tasks did not influence participants' responses. The data were therefore collapsed across these manipulations. A 4 (Shared structure) x 2 (Centrality) ANOVA was performed on the resulting data with repeated measures on both factors. Shared structure had a main effect ($F_1^*(3,21)=4.83$, $p<.05$). The effect of Centrality was significant ($F_1^*(1,23)=14.54$, $p<.005$). The Shared structure x Centrality interaction was also significant, $F_1(3,21)=5.37$, $p<.01$.

The presence of a significant Shared structure by Centrality interaction is surprising in that no such interaction was observed in Experiment 12 (see Table 5.1). Notice that, as in Experiment 12, the centrality effect was lowest in the pure projectibility case, the inference from the observed sample of hippos to all hippos. We speculated therefore that the interaction might be because of the very low effect of centrality for Hippos. To this end, a 3 (Shared structure: rhinos or mice or falcons) x 2 (Centrality: central vs. less central) repeated measures ANOVA was carried out. The main effects of shared structure and centrality were highly significant ($F_1^*(2, 22) = 5.97$, $p<.001$; $F_1(1,22) = 18.19$, $p<.001$ respectively). The shared structure by centrality interaction failed to reach significance ($F_1(2, 22) = 2.75$).

Frequency Task. The results from the Frequency task are summarized in Table 5.6. The mean frequency estimates of the group who completed the frequency task first were almost spot-on the actual estimates. The group who completed the frequency task second, however, overestimated the frequency of the central property.

Table 5. 6 Actual frequency estimates and mean (SE) frequency estimates for each task order for Experiment 13.

	Central Property	Less-Central property
Actual Frequencies	14.00	16.00
Frequency task preceding induction.	14.00 (.98)	15.50 (.70)
Frequency task following induction.	16.83 (1.17)	16.33 (1.22)

As in Experiment 12, preliminary tests showed that the name of the properties did not have any effect on responses. The results were therefore collapsed across this manipulation. A 2 (Order) x 2 (Centrality) ANOVA was performed on the resulting data with repeated measures on the last factor. There were no significant effects: Order ($F_1^*(1,22)=1.83$), Centrality ($F<1$), and Order x Centrality ($F_1^*(1,22)=2.25$).

Although the interaction failed to reach significance, the results are in the predicted direction. When the frequency task came second, the mean estimated frequency for the central property is higher than its counterpart, while the mean estimated frequency for the less-central property was largely unaffected by task order. These results are partially confirmed by one-way analyses of variance at each level of Centrality. For the central property $F_1(1,22)=3.42$, $p=.078$, while for the less central property $F_1(1,22)<1$. Four one-sample t-tests were performed, two for the central and the less-central properties (one for each order), to detect whether each deviated from its respective actual frequency (14 for the central and 16 for the less-central property). The mean judged frequency for the central property when the frequency task came second was significantly higher than its actual frequency, $t(11) = 2.42$, $p<.05$. All the other comparisons yielded non-significant results (for all $t(11)<1$).

2. Familiar properties

Induction Task. Table 5.7 summarizes the results. The results closely resemble those of Experiment 12. Participants were willing to project mud spots to hippos and rhinos but not

to falcons or mice. Their estimates may be influenced by direct experience (having seen many rhinos lots of which had mud spots) or by linking mud spots to some observable property and reasoning on the basis of that. For example, having mud spots is closely associated with living near lakes or swamps. Since rhinos live in such habitats (but not mice or falcons) it is quite likely that they will have mud spots. Participants projected heart with very high confidence to all animal categories. This again validates that the participants understood the task.

Table 5. 7 Mean (SE) inductive strength estimates for the familiar properties of Experiment 13.

Shared structure	<u>Mud Spots</u>	<u>Heart</u>
<u>Hippos</u>	77 (4.45)	99 (.83)
<u>Rhinos</u>	56 (6.86)	99 (.42)
<u>Mice</u>	6 (2.30)	99 (.83)
<u>Falcons</u>	5 (2.23)	99 (.83)

Frequency Task. Table 5.8 summarizes the results for the familiar properties. As in Experiment 12, participants underestimated the frequency of mud spots and heart. We speculated that this may in part be because familiar properties were relatively less salient than central and less-central properties because they were allocated less screen space.

Table 5. 8 Actual and mean (SE) judged frequencies for the familiar properties of Experiment 13.

	<u>Mud Spots</u>	<u>Heart</u>
Actual frequency	17.00	24.00
Judged frequency	15.63 (1.05)	21.00 (.72)

DISCUSSION

The main objective of Experiment 13 was to rule out the interpretation that the results of Experiment 12 were only due to the higher memorability of positive versus negative relations.

To that end, both the central and the less-central features were described 'positively' - three properties depended upon each. The only difference between the central and the less-central features was that the properties that depended upon the central feature were themselves more central. The present experiment using a subtle manipulation of centrality, replicated the main findings of Experiment 12. Across all similarity conditions, central properties were more projectible than less-central ones. Furthermore, participants misjudged the frequency of the central but not of the less-central properties when the induction task was evaluated first, although the statistical analyses were only marginally significant. The results of the present experiments therefore cannot be explained in terms of differential memorability.

A difference between the results of Experiments 12 and 13 was that a significant similarity by centrality interaction was evidenced in the latter but not in the former. We speculated (and supported with evidence) that the interaction of Experiment 13 may be because of a very low centrality effect for the inference to all hippos (a similarly low effect was observed in Experiment 12). This argument is *post-hoc* and therefore needs to be validated by further experimentation.

5.5 EXPERIMENT 14

The results of Experiment 13 are still open to an interpretation that does not appeal to differences in centrality. It is possible that participants preferred to project the 'central' over the 'less-central' feature because the properties that depended upon it were more homogeneous themselves. Animals seem to differ more, for example, in the *amount of green in the eye* rather than in *blood density*. Experiment 14 aimed to rule out this possibility by choosing homogeneous dependent properties for both features. For the less-central feature, the dependent properties are related to hair follicles which is a very homogeneous property of mammals.

METHOD

Participants. A new sample of 52 University of Durham students volunteered to participate.

Design and Procedure. The design and procedure of the present experiment were very similar to those of experiments 12 and 13, with the following exceptions. The names and descriptions of the novel central and less central properties were changed (see sample description below). Their names were not counterbalanced. The less central property (folline) appeared in 16 out of 24 descriptions, whereas the central property (neurine) appeared in either 12 out of 24, or 14 out of 24, or 16 out of 24 descriptions. Moving to the familiar properties, heart was substituted by strong immune system. Both the central (strong immune system) and less central (mud spots) familiar properties appeared in 18 out of 24 descriptions. By keeping their frequency constant, we aimed to detect how centrality affects the perceived frequency and projectibility of familiar properties.

Materials. A sample item that included all properties is presented below (neurine is the central and folline the less-central properties):

This hippo has:

The substance neurine that:

- 1) regulates the level of activity of neurons in the cortex
- 2) controls the structure of neurons in the cortex
- 3) controls the number of neurons per square inch in the cortex

The substance folline that:

- 1) regulates the level of activity of hair follicles in the tail.
- 2) controls the structure of hair follicles in the tail
- 3) controls the number of hair follicles per square inch in the tail

A strong immune system

Mud spots

In the present experiment the property *heart* was substituted by the property *a strong immune system*. Also, the category *falcons* was replaced by the category *bats*, since only mammals have hair follicles. In the current experiment therefore all target animals shared a common superordinate.

RESULTS

1. Central vs. Less-Central: novel properties

Induction task. The results are summarized in Table 5.9. below. At all levels of shared structure, central properties were more projectible than less central properties. The centrality effect was (in absolute terms) smallest when the target category was hippos.

Table 5. 9 Mean (SE) percent inductive strength estimates for the novel properties averaged across frequency conditions of Experiment 14.

Shared structure	Centrality	
	Central property	Less-Central property
<u>Hippos</u>	76.52 (2.35)	74.35 (2.74)
<u>Rhinos</u>	79.56 (2.56)	66.29 (3.79)
<u>Mice</u>	80.38 (3.00)	74.61 (3.84)
<u>Bats</u>	74.00 (3.27)	56.88 (4.91)

Preliminary tests showed that there was no significant effect of order or frequency condition on participants' responses. The data were therefore collapsed across these manipulations. A 4 (Shared structure) x 2 (Centrality) ANOVA was carried out on the resulting data with repeated measures on both factors. The effect of Shared structure was significant; $F_1^*(3,49)=7.30$, $p<.001$. A main effect of Centrality was detected; $F_1^*(1,51)=19.77$, $p<.001$. The Shared structure x Centrality interaction was also significant; $F_1^*(3,49)=5.20$, $p<.005$.

As in Experiment 13, the locus of the interaction seems to lie in the pure projectibility case (the inference from the sample of 24 hippos to all hippos) where no centrality effect was

detected. This is confirmed by a 3 (Shared structure: rhinos or mice or bats) x 2 (Centrality: central or less central property) ANOVA with repeated measures on both factors. The effect of Shared structure and Centrality were highly significant ($F_1^*(1, 50)=9.70, p<.001$; and ($F_1(1, 51)= 19.26, p<.001$). No Shared structure by Centrality interaction was detected ($F_1^*(2,50) = 2.10$).

Frequency task. The results are summarized in Table 5.10. The judged frequency estimates for the central and the less central properties are very close in each task order. Because the actual frequency of the central property was less than that of the less central property, the results show that participants overestimated the frequency of the central property.

Table 5. 10 Actual and mean (SE) frequency estimates for each task order for Experiment 14.

	Central Property	Less-Central property
Actual frequency	12.00, 14.00, or 16.00	16.00
Frequency task preceding induction.	16.36 (.72)	16.56 (.68)
Frequency task following induction.	14.96 (.75)	14.78 (.78)

A 3 (Frequency condition) x 2 (Order) x 2 (Centrality) ANOVA was carried out with repeated measures on the last factor. None of the factors had a significant influence on the results.

2. Central vs. Less-Central: Familiar properties

Induction task. The results are summarized in Table 5.11. The central property was more projectible than the less-central property in all levels of shared structure. The projectibility of both properties was proportional to similarity. The projectibility of the less-central property decreased more steeply with decreasing similarity.

Table 5. 11 Mean (SE) inductive strength estimates for the familiar properties of Experiment 14.

Shared structure	Centrality	
	Central property (Strong immune system)	Less Central property (Mud spots)
<u>Hippos</u>	80.10 (2.19)	76.98 (2.78)
<u>Rhinos</u>	76.52 (2.47)	50.44 (4.30)
<u>Mice</u>	67.42 (3.10)	17.10 (3.40)
<u>Bats</u>	68.73 (3.07)	13.10 (2.90)

A 4 (Shared structure) x 2 (Centrality) ANOVA with repeated measures on both factors was carried out on the data. Shared structure and Centrality had significant main effects ($F_1^*(3,49)=86.05$, $p<.001$; and $F_1(1,51)=113.66$, $p<.001$ respectively). The Shared structure x Centrality interaction was also significant ($F_1^*(3,49)=41.64$, $p<.001$). Notice that, as with the novel properties, the centrality effect was smallest when the target category was hippos.

Frequency task. The results are summarized in Table 5.12. The mean frequency estimates are higher for central than for less-central properties. Further, when the frequency task came second, participants slightly overestimated the frequency of the central property and underestimated the frequency of the less central property.

Table 5. 12 Frequencies and mean (SE) frequency estimates for the familiar properties for Experiment 14.

	Central property (Strong immune system)	Less Central property (Mud spots)
Actual frequencies	18.00	18.00
Frequency task preceding induction.	17.52 (.80)	17.00 (.88)
Frequency task following induction.	18.00 (.75)	16.72 (.81)

A 3 (Frequency condition) x 2 (Order) x 2 (Centrality) ANOVA was carried out with repeated measures on the last factor. Central properties were associated with higher mean judged frequencies; $F_1^*(1,46)=4.78$, $p<.05$. No other main effect, nor second or third order interaction was detected (for all $p>.30$).

DISCUSSION

The main objective of Experiment 14 was to rule out the interpretation that the results of Experiment 13 were due to differences in the homogeneity of the dependent properties. To this end, the properties that depended upon the less-central feature were homogeneous across mammals; they were related to hair follicles that all mammals have. The results from the inductive task replicated those of Experiment 13. At each level of similarity the mean inductive strength estimate of the central property was higher than that of the less-central property. Also, as in Experiments 12 and 13, the centrality effect (in absolute terms) was smallest when the target category was hippos. Critically, the results of the present experiments cannot be accounted for by appeals in the homogeneity of the dependent properties.

For the critical comparison between the central and the less-central novel properties, unlike the previous experiments, the perceived frequency of the central property was not influenced by task order. This might be because the less-central property of the present experiment was relatively more central than the less-central property of Experiments 12 and 13. Having fur, and therefore hair follicles, is a comparatively central feature of mammals since it protects them, for example, from weather conditions. The uniformity and reflectance of eye color is presumably less-central because (for most non biologists like ourselves) almost none of their properties depend on it.

The present experiment also controlled for the actual frequency of the familiar properties; each appeared in 18 out of 24 descriptions. It thus allowed for a meaningful comparison between the central property (strong immune system) and the less-central property (mud spots). For the inductive task, the central familiar property was more

projectible than the less central familiar property for each similarity condition. For the frequency task, the mean estimate for the central familiar property was slightly higher when the frequency task came second compared to first, whereas the mean estimate for the less central property were slightly lower for the same task order. Although in the expected direction, the Task order by Centrality interaction failed to reach significance. However, these results closely resemble the frequency judgment data for the novel features of Experiments 12 and 13. Put together, the data tentatively suggest that centrality can bias the perception of frequency.

A surprising finding was that the frequency condition influenced neither the frequency nor the inductive strength estimates. It seems therefore that participants were largely insensitive to observed frequencies.

5.6 GENERAL DISCUSSION

The present experiments had two main objectives: (1) to examine the relative contributions of centrality and variability information in judgments of projectibility and frequency (all experiments), and (2) to rule out the interpretation that the effect of centrality on projectibility reported in Chapters III and IV can be explained solely in terms of feature overlap with respect to the dependent properties. The present findings are discussed in terms of each issue.

5.6.1 Centrality versus frequency contributions on projectibility

In the present experiments, participants were given a choice between conceptual structure information (centrality) and statistical information (frequency) to make an inductive response. The results showed that inductive responses were based on centrality information and were largely unaffected by observed frequencies.

Further evidence that information about conceptual structure is preferred to frequency information comes from Murphy and Allopenna's (1994) studies in category learning and use.

Participants who were presented with categories whose features could be related by a common theme (schema), such as "arctic vehicle" or "jungle vehicle," learned the categories faster than participants who received features unrelated to a common theme (note that even in the common theme group the 'theme' was left implicit). The former participants were as fast and as accurate to identify any property of a category as a typical one, irrespective of its actual frequency. Importantly, they rated both frequent and infrequent properties as about equally typical. As in our study, it seems that the thematic participants based their responses not on the actual frequencies of properties but rather on how well the property hanged together in the structure of the thematic category.

Broadly, our results echo Murphy and Kaplan's (1997; p. 170) empirical claim that "when subjects have the choice of using knowledge or the statistical structure of the category in making a response, the knowledge tends to win" (see e.g., Ahn et al., 1995; Murphy & Kaplan, 1997; Spalding & Ross, 1994; Wisniewski, 1995).

5.6.2 Centrality information may bias the perception of frequency

Centrality information was found to bias the perception of frequency for the novel properties in Experiments 12 and 13, although only marginally in Experiment 13. Participants who estimated the induction task first overestimated the frequency of the central property. This tendency was not observed in Experiment 14. We stipulated that this might have been because the 'less-central' property (which was related to *having fur*) was relatively central. A tendency to overestimate the frequency of the central with respect to the less-central property was also obtained for the familiar properties of Experiment 14 (strong immune system versus mud spots). All in all, we take the data to suggest that centrality information may bias the perception of frequency.

5.6.3 Centrality and projectibility

A possible joint interpretation of the results of Chapters III and IV is that features are projectible to target concepts in proportion to the extent that the target concepts share their

dependent properties. Following Sloman et al. (1998) we assume that centrality can be defined recursively: a feature's centrality depends upon the features' degree of centrality that depend upon it, which in turn depend upon the feature centrality of the properties that depend upon them and so forth. Features therefore upon which more central properties depend are more central than features upon which less central properties depend. According to our hypothesis, therefore, we expected the former features to be more projectible than the latter. In addition, because all the dependent properties were general, the degree of overlap between the base and target should be constant across similarity conditions. Hence, to the extent that the effect of centrality is proportional to the degree to which the target shares the depending on properties specified for the base, we expected no shared structure by centrality interaction. Both of these predictions were confirmed: central properties were more projectible than less central ones (bar the pure projectibility case) and the centrality effect was largely independent of shared structure level.

The present experiments also examined, for the first time, pure projectibility cases (known as generalization from instances); e.g. inferences from a sample of hippos to all known hippos. The results suggest that central properties are more projectible than less central ones, though this effect was of a smaller magnitude than that of other shared structure conditions. In retrospect, this result makes sense. Hippos share with one another lots of characteristics and properties lots of which are mutable (e.g. color, smell, mud spots, stains in the trunks and so forth). In this case homogeneity information seems more relevant in determining the extent to which all hippos have a property: there are lots of properties that hippos share but only for few of them we have theories that constitute them immutable. Further, both the central and the less central properties were intrinsic (e.g. hormone, substance) and therefore had some inductive potential to begin with. Possibly one could observe a significant effect of centrality on the pure projectibility case by: (1) giving information from only a small sample of a category (say two hippos) and / or by (2) using intuitively mutable novel features (like "shiny teeth") and for the centrality condition stipulate that some very central properties depend upon them (e.g. reproductive success).

CHAPTER VI:
CENTRALITY VERSUS VARIABILITY:
INFORMATION SEARCH

6.1 CONCEPTUAL CENTRALITY VERSUS VARIABILITY: INFORMATION SEEKING

The series of experiments reported in Chapter V showed that people, when provided with both frequency and centrality information, they use centrality information rather than frequency information to draw inferences. Reviewing evidence on information use, we concluded that when people are given a choice between statistical or causal mechanism information to make a response they base their response on the latter. Taking a step back, one could ask what sort of information would people *seek* to make a response. This is the main question of the present chapter: Do people seek mutability or variability information to make an inference?

Suggestive evidence that people seek information about causal mechanisms rather than statistical information comes from the domain of causal attribution. Traditional models of causal attribution suppose that people in order to determine the cause of an event seek information about the co-variation of factors (e.g. Kelley, 1967, 1973; Cheng & Novick, 1990, 1992). However, Ahn, Kalish, Medin, and Gelman (1995) provided evidence on information seeking strategies that challenges this view. Specifically, they found that to make a causal attribution people ask more about underlying causal mechanisms than about co-occurring events. For instance, when trying to discover why John had an accident on route 9 last night, participants tended to ask questions that presupposed causal mechanisms (e.g. was John drunk?), rather than questions about the frequency that John had accidents on other occasions, or the frequency of other car accidents last night like the co-variation models presuppose.

Covariation models of causal attribution with their focus on common versus distinctive factors are closely related to similarity-based models of category-based inference with their emphasis on common versus distinctive features (such as Sloman's (1993) feature-based model). Recall that similarity-based models work quite well with blank properties - properties about which participants have very little knowledge. Centrality information presupposes some mechanism ('intuitive theory') that binds the features of a concept together.

Even vague centrality information (e.g. lots of a lion's features depend on the neurotransmitter dihedron) gives some clue about the causal status of the candidate feature. To the extent therefore that people seek information about mechanisms to make a response, we expected them to seek centrality over raw frequency information to make an inference.

The weight of centrality versus frequency information may critically depend on the similarity between the base and the target category. Similar animal categories (like Indian elephants and African elephants) share many properties: almost all of their central ones (which are relatively few) but also many non-central ones (e.g. color of eyes, hardness of nails, and size). Dissimilar animal categories (like zebras and chimps) share some central properties and very few mutable properties. Finding, for instance, that zebras have a property upon which many (or very few) of its functions depend, presumably gives some reason to believe that chimps will have (or will not have) that property. Knowing that many (or very few) zebras have a property provides no compelling evidence that chimps will (or will not) have that property. In other words, it is possible that the relative value of centrality and homogeneity information depends on the degree of similarity between the premise and conclusion categories. Specifically, the usefulness of getting centrality over frequency information seems to increase as the similarity between the base and the target categories decreases.

The present chapter presents two experiments that aimed to detect what type of information people seek (centrality or variability) to make a categorical inference. The participants' task was to decide about whether a target animal had a novel candidate property. They were given a choice between variability and centrality information. To detect whether the preference to seek centrality over variability information varies depending on the level of shared structure, the similarity between the premise and conclusion categories was varied. To the extent that people prefer information about causal mechanisms rather than raw frequencies, we expected them to seek centrality over variability information. Also, to the degree that this preference is mediated by the similarity between the base and the target

categories, we expected the preference to choose centrality over variability information to be proportional to the similarity of the animal pair.

6.2 EXPERIMENT 15: CENTRALITY VS VARIABILITY IN INFORMATION-SEEKING

METHOD

Participants. Twenty-four first-year undergraduates in psychology from the University of Durham voluntarily participated in this study.

Design and Materials. Participants had to choose between centrality and variability information about a property for an animal (base) in order to make a choice about a different animal (target). The similarity between animal-pairs was varied (high-similarity same-superordinate vs. low-similarity same-superordinate, vs. low-similarity different-superordinate) in a repeated measures design.

The present experiment used the same animal pairs and types of properties (enzymes, hormones, or neurotransmitters) as Experiment 6. There were six base-target pairs, each of which was modified to form 6 triples, with one member of the triple in the HSSS condition, one in the LSSS condition, and one in the LSDS condition. The only difference between members of a triple was the category used as a premise. The base categories {lions, raccoons, sparrows}, for instance, were paired with the target category {tigers} to form a triple. The full set of premise triples is shown in Table 3.7 and the full set of properties in Table 3.9.

In each premise, two pieces of information have been blanked out, one before the category name (Blank 1) and one before the last phrase in the premise concerning the extent to which relevant functions depend on the cited properties (Blank 2). Blank 1 represents variability information whereas Blank 2 represents centrality information. After the inference task has been described, participants were instructed to say which of the two blanks they would like to have filled-in to complete the task.

As in Experiment 6, three different lists were constructed that counterbalanced the assignment of type of property to premise triples (see Table 3.8).

Procedure. Participants were given a booklet containing one of the three lists of materials. Participants were asked to imagine that they were finalists in a TV quiz show, 18 questions away from winning the grand prize. Their final task was to estimate the likelihood that members of an animal species have a physiological property. They were told, "The presenter gives you a choice of one of two pieces of information, both of which concern another species of animals." Then, participants were given an example similar to the test items to complete. They were asked to choose the Blank that they would rather have filled-in to complete the task, and that beforehand they had no clue about the precise information that Blanks 1 and 2 contain.

RESULTS

Coding. The choice of centrality information was coded as 1 and the choice of variability information as 0. The dependent variable was calculated for each participant as the number of Centrality choices (number of 1s) in each Similarity condition.

The results are summarized in Table 6.1. Participants chose centrality over variability information as the similarity of the animal pairs decreased.

Table 6. 1 Mean choice of Centrality information in each similarity condition.
Chance level = 3.00.

Shared structure	Mean (SE)
<u>HSSS</u>	2.54 (.45)
<u>LSSS</u>	4.29 (.46)
<u>LSDS</u>	4.92 (.36)

Note. HSSS = High-Similarity Same-Superordinate.
LSSS = Low-Similarity Same-Superordinate.
LSDS = Low-Similarity Different-Superordinate.

Shared structure. Two 3 (Shared structure) x 3 (List) analyses of variance were conducted. In the F_1 analyses Shared structure was a repeated measures factor, in the F_2 analyses both factors were repeated measures. The main effect of Shared structure was significant ($F_1^*(2,20) = 9.56, p < .005$; $F_2(2,15) = 80.83, p < .001$). No main effect of List was found ($F_1(2,21) = 1.73$), nor a significant Shared structure x List interaction ($F_1 < 1$). Moving from the HSSS to the LSSS condition resulted in a significant 1.75 point increase (participants: $t(23) = -3.17, p < .005$; items: $F_2(1,10) = 67.85, p < .001$). Moving from the HSSS to the LSDS condition resulted in a significant increase of 2.38 points (participants: $t(23) = -4.31, p < .001$; items: $F_2(1,10) = 122.14, p < .001$). The means of the LSSS and the LSDS conditions were marginally different for participants ($t(23) = -1.97$), but significantly different for items ($F_2(1,10) = 18.75, p < .005$).

Centrality or Variability? Participants in the low- but not in the high-similarity conditions sought more centrality information. Two-tailed one sample t-tests (chance level = 3) for participants at each level of Similarity confirm this. The means for the LSSS and LSDS conditions were significantly higher than chance; $t(23) = 2.83, p < .01$ and $t(23) = 5.39, p < .001$, respectively. The mean for the HSSS condition was lower than chance but not significantly ($t(23) = -1.09$). The result that the centrality choice in the LSSS and LSDS conditions was significantly greater than chance was replicated by one-sample t-tests across items ($t(5) = 12.87, p < .001$, and $t(5) = 18.18, p < .001$ respectively; chance level = 12).

DISCUSSION

Participants preferred centrality to variability information to make a categorical inference among low similarity animal pairs. For highly similar animals no such preference was observed. In general, the preference to seek centrality over variability information was negatively related to the judged premise-conclusion similarity.

6.3 EXPERIMENT 16: A REPLICATION OF EXPERIMENT 15

Experiment 15 had a shortcoming: its instructions were ambiguous. The instructions failed to specify the precise sort of information that was represented by the blanks. The present experiment aimed to replicate the results of Experiment 15 while making crystal clear that the blanks would filled-in with vague centrality ("few/lots" type) and variability information.

METHOD

Participants. Thirty-one first-year undergraduates in psychology from the University of Durham voluntarily participated in this experiment.

Design and Materials. The design of the present experiment was the same as that of Experiment 15, except of the wording of the items. Participants were given a choice between centrality and variability information about a property of an animal to make a categorical inference about another animal. The similarity between animal-pairs was varied (high-similarity same-superordinate vs. low-similarity same-superordinate, vs. low-similarity different-superordinate) in a repeated measures design. The six triples used in Experiment 15 were used in this experiment (see Table 3.7), but the wording of the items was different. Examples using a triple are shown in Table 6.2 overleaf. For the full list of properties used see Table A.6.3.

Two lists of materials were constructed to counterbalance the presentation order of the centrality and variability information. In List 1 for triples 1, 3, and 5 centrality information was presented first, whereas for triples 2, 4, and 6 centrality information was presented second. In List 2 this presentation order was reversed.

Table 6. 2 Sample items from Experiment 15.

Shared structure

HSSS Task: You want to find out whether chimps have the substance trolone.

Circle the question below that you would most like to have answered to help you with the task.

- A. How many gorillas have trolone?
- B. For the gorillas that have trolone, how many of their properties and functions depend on that substance?

LSSS Task: You want to find out whether chimps have the substance panctone.

Circle the question below that you would most like to have answered to help you with the task.

- A. How many zebras have panctone?
- B. For the zebras that have panctone, how many of their properties and functions depend on that substance?

LSDS Task: You want to find out whether chimps have the substance zylone.

Circle the question below that you would most like to have answered to help you with the task.

- A. How many blackbirds have zylone?
 - B. For the blackbirds that have zylone, how many of their properties and functions depend on that substance?
-

Procedure. Participants received a booklet containing one of the two lists of materials.

Participants read the following instructions:

Thank you for taking part in this study! Your task is to find out whether members of an animal species have a particular substance. To help you with the task you can ask one and only one question concerning the members of another animal species. Consider an example:

Task: You want to find out whether cats have the substance zioline.
Circle the question below that you would most like to have answered to help you with the task.

- A. How many dogs have zioline?
- B. For the dogs that have zioline, how many of their properties and functions depend on that substance?

Possible answers to these questions are:

Almost all Very many Many Few Very few Almost no(ne)

In the case you have any queries, please ask the experimenter now. Otherwise turn the page and begin.

RESULTS

Coding. The choice of centrality information was coded as 1 and the choice of variability information as 0. The dependent variable was calculated for each participant as the number of centrality choices (number of 1s) in each shared structure condition.

The results are summarized in Table 6.3. Participants show a preference of centrality over variability information in all similarity conditions. This preference is slightly more evident in the low rather than in the high similarity conditions.

Table 6. 3 Mean choice of centrality information in each shared structure condition. Chance level = 3.00.

Shared structure	Mean (SE)
<u>HSSS</u>	4.03 (.40)
<u>LSSS</u>	4.90 (.40)
<u>LSDS</u>	4.74 (.39)

Shared structure. A preliminary analysis showed that the type of list did not have an effect on the participants' responses. The results therefore were collapsed across lists. The effect of Shared structure was not reliable. It was not significant across participants ($F_1^*(2,29) = 1.85$) but it was across items ($F_2(2,15) = 23.81, p < .001$).

More centrality information was sought in the low- rather than in the high-similarity conditions. Moving from the HSSS to the LSSS condition resulted in a .87 point increase (participants: $t(30) = -1.72$; items: $F_2(1,10) = 37.58, p < .001$). Moving from the HSSS to the LSDS condition resulted in an increase of .71 points (participants: $t(30) = -1.47$; items: $F_2(1,10) = 122.14, p < .001$). The means of the LSSS and the LSDS conditions were not different (participants: $t(30) = -1.41$, items: $F_2(1,10) = 1.92$).

Centrality or Variability? In each similarity condition centrality information was sought significantly more times than variability information. The one-sample t-statistics were: HSSS condition (participants: $t(30) = 2.53, p < .05$; items: $t(5) = 8.88, p < .01$), LSSS condition (participants: $t(30) = 4.82, p < .001$; items: $t(5) = 23.31, p < .001$), and LSDS condition

(participants: $t(30) = 4.37, p < .001$; items $t(5) = 21.01, p < .001$). The chance level for participants was 3, for items 15.5.

DISCUSSION

Experiment 16 aimed to make it crystal clear that the centrality and variability information that participants would get would be vague (e.g. for the hippos that have the candidate feature, lots of properties depend on it; lots of hippos have the candidate feature). As in Experiment 15 participants showed an overall preference for centrality information, but this time in all shared structure conditions. Further, similarly to Experiment 15, there was some suggestion that the preference to seek centrality over variability information was higher for low similarity rather than for high similarity animal pairs. However, the trend was not significant in this experiment.

6.4 GENERAL DISCUSSION

The present chapter reported two experiments that aimed to determine whether people prefer centrality to variability information to make an inference, and, in the case they do, whether the amount of shared structure between the premise and conclusion categories interacted with this preference. Taken together, experiments 15 and 16 suggest that participants overall seek centrality over variability information to make an inference. This preference is especially evident for inferences involving low similarity animal pairs.

In the introduction we noticed that in the domain of causal attribution recent research has shown that people tend to seek information about causal mechanisms rather than about co-variation (see Ahn et al., 1995) to decide the cause of an event. Similarly, the present experiments found that people generally seek centrality information (information about how much of the internal structure of a concept depends on that feature) over raw frequency information to decide whether a target category had a novel property. Taken together, these studies suggest that when people seek information to make a response they prefer information

about the underlying mechanisms (or the role of a feature or factor in this mechanism) rather than raw statistical information.

CHAPTER VII:
EXTENDING THE CENTRALITY
HYPOTHESIS TO ARTIFACTS

7.1 NATURAL KINDS VERSUS ARTIFACTS

The present account assumes a domain-general inference mechanism as do most psychological models of category-based inference (e.g. Osherson et al., 1990, 1991; Sloman, 1993). At the same time, most of such models have only been tested with natural kind categories (for exceptions see Sloman, 1998). It follows that claims for all-purpose inference mechanisms should be tested outside the domain of natural kinds. This need is even more pressing given that many recent studies have demonstrated effects of domain specific knowledge in concept learning and use (e.g. Carey, 1985; Murphy & Medin, 1985; Pazzani, 1991; Wattenmaker, 1995; for a review see Heit, 1997). Such evidence appears to challenge the possibility that a domain general inference process underlies category-based induction (cf. Hirschfield & Gelman, 1994). The present chapter begins with a review and evaluation of conceptual and empirical evidence on differences between artifacts and natural kinds. Subsequently it presents an experiment that aims to extend the centrality hypothesis to artifacts.

7.1.1 Conceptual evidence and evaluation

Numerous factors have been suggested to contrast natural kind categories (such as animals, gold, and tornadoes) and artifact categories. To name a few, artifact categories are said to differ from natural kind categories in that: (1) they do not occur naturally; (2) they do not have essences (whatever these turn out to be); (3) they do not have sciences centered around them; (4) they require reference to human intention.

These distinctions have not been left unchallenged. Natural kinds are not the only naturally occurring classes of things because some other classes of things are also naturally occurring. Blue things, for instance, describe a naturally occurring class. Further, not all natural kinds may yet exist. Some classes of things, like a particular radioactive element, may be sufficiently rare as to not have been observed yet (see Keil, 1995). Naturalness fails therefore to segregate natural kinds from artifacts.

Natural kinds are not even the sole classes of things that have sciences centered around them - artifacts are also being studied (e.g. antique experts study antiques, dress designers study fashion). Gelman (1988) suggests that a critical difference between artifact experts and natural kind experts might be that the former focus on complex sets of distinctions whereas the latter on general laws that describe entire categories (such as mammals). This claim is based on the assumption that artifacts are inherently more variable than natural kinds. Almost all features of artifacts can vary as long as their intended function is satisfied (e.g. telephones can vary in color, shape of handle, and material as long as they can function as telephones). In contrast, natural kinds are constrained by their genetic make-up to be of a certain color, to reach a certain size, and so forth. Notice, however, that the heterogeneity of properties varies within a domain. For example, color is a homogeneous property for elephants and PCs but not for parrots and shirts. Inductive potential should not therefore depend on category kind, but rather on the relation between a predicate and a category (Goodman, 1955; Markman, 1989; Nisbett, Krantz, Jepson, & Kunda, 1983).

Reference to human intention also seems inadequate to delimit artifacts from natural kinds. Natural kinds, such as designer plants or animals subject to intensive breeding, seem to require reference to human intention. Vice versa, complex artifacts, such as computers, do not seem to require reference to human intention (at least to a certain extent).

Based on such arguments theorists like Keil (e.g. 1989, 1995) contend that natural kinds and artifacts are arrayed along several continua rather than in contrasting bins. This leaves the door open for general mechanisms, such as conceptual centrality, to capture category-based inference across various domains.

7.1.2 Empirical evidence and evaluation: Type of property by type of kind interaction.

One sort of empirical evidence in support of the domain-specific approach comes from studies that show a type of property by type of category interaction (e.g. Gelman 1988; Barton & Komatsu, 1989; Keil; 1989, 1995; Rips 1989). These studies show that internal or

molecular features are important for categorizing natural kinds (e.g. for bird: it mates with other birds), whereas external, functional features are important for categorizing artifacts (e.g. for coffeemaker: it is used for making coffee). One interpretation of such findings is that natural kinds and artifacts involve different domains and hence the interaction between property and category type is an immediate consequence of this domain distinction. A review of the literature suggests that this claim is unwarranted.

Gelman (1988, Study 1) studied inductive inferences about natural kind and artifact categories with four-year olds and second-graders. One of her findings was that second graders were more willing to generalize among artifacts' functional features (e.g. for clock: used for horology; for bike: you can traverse with it), rather than molecular or previous form features (e.g. for clock: has a pondus inside; for bike: used to be a piece of bauxite).

Similar evidence is reported from the domain of categorization. Keil (1989), for example, had children consider an instance of a category lacking some surface properties and asked them whether they still considered the transformed instance as a member of the original category. By systematically varying the type of category, Keil found that surface features were rated more important for the identity of artifacts than for the identity of natural kinds. One story, for instance, involved a raccoon that underwent a transformation and came to look like a skunk (e.g. it smelled bad, and its tail was painted black with a white stripe). Most of the older children judged the transformed animal to have kept its identity, to still be a raccoon. Another story involved the transformation of a coffeepot into a birdfeeder. Most of the older children judged the transformed artifact to have lost its identity, to be a birdfeeder. (For similar evidence see also Rips, 1989; Barr & Caplan, 1987; Barton & Komatsu, 1989).

Contrary to these studies, Malt and Johnson (1992) presented evidence against the idea that functional features are the most important in category decisions for artifacts. They found that some physical features (e.g. for tractor: has an engine, has large wheels, has one seat, the seat is unenclosed, has an attachment for farm machinery) are more important than, or at least as important as, some functional features (e.g. for tractor: allows one person on a

farm to till ground or plow fields by pulling a variety of other machines). Malt and Johnson's findings cast therefore doubt to the validity of the property type x category type distinction.

In sum, both the conceptual and empirical evidence fail to support a clear-cut distinction between artifacts and natural kinds. At the same time, the bulk of evidence suggests that there are broad differences between these types of categories. These broad differences seem to cast doubt as to whether a domain general mechanism can account for category-based inference. Below we present evidence that conceptual centrality offers just such a domain general mechanism.

7.2 CONCEPTUAL CENTRALITY CAN ACCOUNT FOR CATEGORY BY PROPERTY INTERACTION

Ahn's (1998) studies on adult categorization provide evidence that the causal status hypothesis (basically that causes are weighted more heavily than effects in categorical decisions; see Chapter IV for a more detailed discussion) can account for why functional features (e.g. it is used to pound nails) are generally more important for the identity of artifacts, whereas molecular features (e.g. it has a left aortic arch) for the identity of natural kinds. Her first two studies borrowed the materials from Barton and Komatsu (1989) and Malt and Johnson (1992) respectively, and showed that in both cases conceptual centrality was correlated with judged category importance. The results therefore give some credibility to the claim that causal centrality constrains real-life categorization. In her last two studies, Ahn examined artificial natural kind and artifact categories while directly manipulating the causal status of a property. She found that causal status was a good predictor of category importance independent of the type of category involved. Functional or molecular high-causal status features were judged as more important in categorical decisions than functional or molecular low-causal status features across both artifacts and natural kind categories. Taken together, Ahn's experiments strongly suggest that a feature's causal status constrains categorization.

Since causal status is but a special case of conceptual centrality, her studies corroborate the idea that conceptual centrality constrains category-based inference.

We hold that artifacts are similar to natural kinds in the sense that they are also embedded in lay theories (or vague pre-theoretical biases), theories that can be represented by asymmetric dependency links connecting their properties. Otherwise stated, we hold that the features of artifacts also differ in terms of conceptual centrality and therefore they should be differentially projectible. *Being able to freeze-up* stuff, for instance, is a conceptually central feature of refrigerators, *being white* is not. Much like with animal categories, we expect, all else being equal, that the more central a feature for a concept the higher its projectibility to other concepts that share its attributes and dependencies.

7.3 EXPERIMENT 17: IT'S ALL GREEK TO ME

Our experiments so far have tried to minimize participants' background knowledge of the central and less-central candidate features in an effort to keep 'all else equal'. Animal categories, being highly complex, are well suited to study our hypothesis in part because they possess many properties that we simply do not know about. Few people outside biological sciences, for instance, know the names of the various hormones that animals have. For artifacts, and especially for simple ones like chair, it is hard to come up with credible novel properties that are unbeknown to the participants¹⁰. Because of this limitation, in the present experiment the candidate properties were named in Greek. Relative centrality was defined from the "Few/Lots" definition. Participants were instructed that they would learn two facts about a category, and that they would be asked to generalize these to another category. To motivate them we told them that this task was in a context of a TV-show that promised a large sum of money for the best answers.

The main objective of Experiment 17 was to show that feature centrality influences property induction among artifacts. Our straightforward prediction was that central features

would be more projectible than less-central features to the extent that the argument's categories shared structure. To investigate what sort of structure the categories should share we included 3 levels of similarity (high-surface high-functional (HSHF), or high-surface low-functional (HSLF), or low-surface high-functional (LSHF). According to our prediction the centrality effect should be most evident in the HSHF level. Also, to the extent that shared functions mediate more the centrality effect, we expected a stronger centrality effect in the LSHF than in the HSLF condition. Such a result would sit well with theories supporting functional features as central for artifacts. Such a result, however, is independently expected because centrality is defined in terms of how many functions of an artifact depend on it. That is, LSHF category pairs are more similar in the relevant dimensions than HSLF pairs.

METHOD

Participants. The participants were 22 first-year undergraduates of the University of Durham. Participation was in the context of a tutorial.

Design. This experiment crossed Centrality (central or less-central property) with Shared structure (high-surface high-functional (HSHF), or high-surface low-functional (HSLF), or low-surface high-functional (LSHF)) with repeated measures on both factors.

Materials. Eight triples of materials were constructed in the same way as the triples in Experiment 15. That is, one member of a triple appeared in the HSSS condition, one in the LSSS condition, and a third in the LSDS condition. The only distinguishing feature between conditions was the critical category in the premise of the argument. An example triple is shown in Table 7.1. As can be seen from the table, the wording of the items is the same as the one used in the few/lots set of experiments (Chapter IV).

¹⁰ The claim is not that complexity differentiates natural kinds from artifacts. Complexity varies within a domain; e.g. water is less complex than zebra; pencil is less complex than microwave oven.

Selection of base-target pairs. A series of base-target pairs were constructed by the author that were thought, intuitively, to differ in their surface similarity and function similarity. They were then discussed with his supervisor and a number of revisions, deletions and additions were made. The list was then further refined by the author to produce a list of 24 items, eight in each Shared structure condition. This list was then presented to an independent sample of 12 participants who were asked to rate each item for its functional and surface similarity. For the functional similarity estimates participants were instructed (the actual instructions were in bold-faced letters): "Please rate how similar are the following pairs of objects in terms of function. The higher the rating the more functionally similar you think that the pair of objects are." For the surface similarity estimates participants were instructed: "Please rate how similar are the following pairs of objects in terms of their surface properties (e.g., how do they look). The higher the rating the more similarly looking you think that the pair of objects are."

Items were constructed in a similar way as in Experiments 6, 7, 10, and 11, with the exception that in the present experiment a single property was assigned to each base-target pair. Since the properties were in Greek we felt it unnecessary to counterbalance the assignment of properties to premise triples. Eight premise triples were paired with a single target category to form an argument in each shared structure condition. For example, the premise triple {bungalow, Barbie's house, tent} was paired with {house}. Table 7.1 overleaf presents the 24 argument categories of Experiment 17.

In each questionnaire half of the questions in each shared structure condition asked for the likelihood of the target having the central property first, and the other half second. The order of which the central and the less central properties were evaluated first, was counterbalanced across participants. The presentation order of the items was also counterbalanced such that about half the participants received one order while the rest its reverse.

Table 7. 1 Premise triples along with the target categories for the items of Experiment 17.

Premise Triples			Target
<u>HSHF</u>	<u>HSLF</u>	<u>LSHF</u>	
Lorry	Toy-truck	Cargo-boat	Truck
Trumpet	Plastic saxophone	Music-box	Saxophone
Washing-machine	Refrigerator	Car-wash	Dish-washer
Clock	Compass	Hour-glass	Watch
Mac	TV	Calculator	IBM-PC
Microwave	Freezer	Camp-stove	Oven
Boeing 747	Remote-control plane	Rocketship	Concorde
Bungalow	Barbie's house	Tent	House

Procedure

Similarity task. The items were presented in booklets. For the functional similarity estimates participants were instructed (the actual instructions were in bold-faced letters): "Please rate how similar are the following pairs of objects in terms of function. The higher the rating the more functionally similar you think that the pair of objects are." For the surface similarity estimates participants were instructed: "Please rate how similar are the following pairs of objects in terms of their surface properties (e.g., how do they look). The higher the rating the more similarly looking you think that the pair of objects are."

Induction task. The items were presented in booklets. Participants were asked to imagine that they were finalists in a TV quiz show few questions away from winning the grand prize: £100,000. The presenter unveiled the last task named "It's all Greek to me!" They were told that they would be informed that an object had two properties. Their task was to rate the likelihood of another object having each of these properties. The catch was that the names of the properties were given in Greek. Participants had to work through an example before proceeding with the test items.

RESULTS

Similarity results. Table 7.2 summarizes the results from the similarity control experiment. Broadly, the original assignment of category pairs to similarity conditions fits well with the participants' judgments. HSHF items were, as intended, high in both surface and functional similarity. HSLF items were higher in surface similarity but lower in functional similar than LSHF items.

Table 7. 2 Mean (SE) functional similarity and surface similarity ratings for the 3 similarity conditions. The higher the rating, the higher the judged similarity.

Similarity	Similarity type	
	Functional Similarity	Surface Similarity
<u>HSHF</u>	77.72 (5.97)	67.37 (4.67)
<u>HSLF</u>	17.55 (2.65)	55.97 (4.60)
<u>LSHF</u>	43.66 (4.55)	20.90 (5.08)

Note. H=High; L=Low; S=Surface; F=Functional.

To detect whether the assignment of items to similarity condition were justified two 3 (Similarity: HSHF or HSLF or HFLS) x 2 (Similarity Type: functional or surface) repeated measures analyses of variance were performed, one across participants and one across items. The only robust effects were that of Similarity and the Similarity x Similarity Type interaction. For Similarity the F-values were: $F_1(2,21)=41.11$, $p<.001$; $F_2(2, 10)=11.12$, $p<.005$. For the interaction the F-values were: $F_1(2,21)=25.69$, $p<.001$; $F_2(2,10)=40.99$, $p<.001$. These results sit well with the assignment of properties into conditions. The effect of similarity is because the HSHF condition is high in both types of similarity whereas the HSLF and HFLS in just one. The interaction is because of the differences between the HSLF and HFLS conditions.

To test this latter proposal we performed a 2 (Similarity: HSLF or HFLS) x 2 (Similarity type: functional or surface) repeated measures analyses of variance. As expected,

the only robust effect was the Similarity x Similarity type interaction ($F_1(1,14)=48.05$, $p<.001$; $F_2(1,11)=50.71$, $p<.001$). The assignment of items into the HSHF, HSLF and LSHF conditions seem therefore to be justified.

Induction results. Table 7.3 summarizes the results. For the HSHF condition a strong centrality effect is observed. For the LSHF and the HSLF conditions central properties were more projectible than less central ones, though the value of the difference is much lower. The centrality effect for the LSHF condition is slightly higher than that of the HSLF condition.

Table 7.3 Mean (SE) inductive strength estimates for Centrality by Similarity conditions.

Similarity	Centrality	
	Central	Less-Central
<u>HSHF</u>	61.91 (2.28)	42.88 (3.42)
<u>HSLF</u>	33.72 (2.80)	30.76 (3.11)
<u>LSHF</u>	35.37 (2.66)	30.71 (3.30)

The counterbalancing methods did not have a significant effect on the responses, and hence they were dropped from the analyses. Two 2 (Centrality) x 3 (Similarity) analyses of variance were carried out. In the F_1 analysis both factors were repeated measures, in the F_2 only Centrality was repeated measures. There was a main effect of Centrality ($F_1^*(1,21)=12.91$, $p<.005$; $F_2(1,21)=34.83$, $p<.001$), and a main effect of Similarity ($F_1^*(2,20)=58.80$, $p<.001$; $F_2(2,21)=17.56$, $p<.001$). The interaction was also significant ($F_1^*(2,20)=8.48$, $p<.005$; $F_2(2,21)=11.43$, $p<.001$).

To find the locus of the interaction, t-tests were performed across each level of similarity. A significant centrality effect was shown for the HSHF condition (participants: $t(21)=4.75$, $p<.001$; items: $t(7)=33.88$, $p<.001$). For the LSHF condition, the results were significant only for participants (participants: $t(21)=2.09$, $p<.05$; items: $t(7)=3.84$). For the HSLF similarity condition, the results did not reach significance (for both $t<1$).

DISCUSSION

The aim of this experiment was to extend our conceptual centrality hypothesis to the artifact domain. The results clearly support this claim. A secondary aim was to detect whether the centrality effect would be mediated more heavily by functional rather than by surface similarity. Although the results somewhat suggest that the centrality effect is more closely associated with functional similarity (a centrality effect was obtained for the LSHF but not for the HSLF condition), they only do so mildly. The only clear cut finding is that centrality *does* influence property induction with artifacts, and especially among concepts that they are highly similar on both surface features and deeper functions. The main effect of centrality is strongly modified by the centrality by similarity interaction.

Although this experiment evidenced an influence of centrality on category-based inference with artifacts, it is not clear what drove participants' judgments. At one end, participants may have reinterpreted the Greek names into known properties of the premise categories. If this reinterpretation was random, then both the central and the less-central properties should have been equally projectible. Since a centrality effect was detected, this possibility is discarded. At the other end, participants may have used a general centrality heuristic, something like "project the central property no matter what". This second possibility must also be discarded since it fails to explain the overwhelming centrality x similarity interaction. A third possibility is that people reinterpreted the properties into known properties that fitted the bill, which agreed with the centrality descriptions. For instance, they might have reinterpreted the central property "mixani" of lorry as "engine" (which by the way is its real meaning), and the less-central property "kathisma" as "seat". An account of the results therefore is that central novel properties were reinterpreted as more generalizable known properties. Notice though that the only information participants were supplied with was vague information about a feature's centrality. So, the possibility that central novel features cue generalizable familiar features supports that feature centrality influences generalizability.

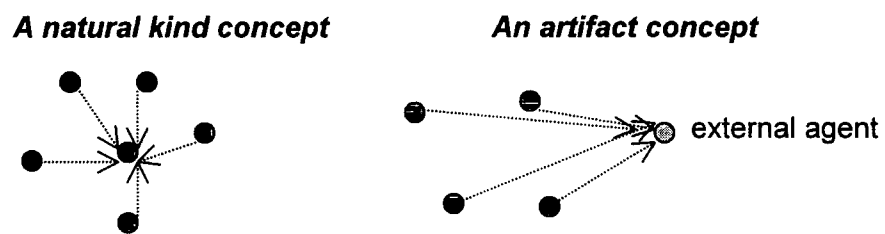
To summarize, the results show that centrality *does* exert an influence on property induction with artifacts in proportion to the similarity between the base and the target

concepts. Needless to state, a complete account of how people reason about artifact categories requires further experimentation.

7.4 GENERAL DISCUSSION

Chapters II to VI investigated the influence of centrality on property induction with categories from the natural kind domain. The present chapter extended this hypothesis to the artifact domain. Central properties were more projectible than less-central ones in proportion to the similarity between the base and the target concepts. Taken together, the findings suggest that our centrality hypothesis can account for category-based inference across different domains (see also Ahn, 1998).

Figure 7. 1 A simplified illustration of the view that dependencies point inwards for natural kinds but outwards for artifacts. Dots represent properties. (adapted from Keil, 1995).



To avoid misinterpretations, we do not doubt that there are differences between artifact and natural kind concepts. Such differences, however, may involve more how dependency patterns in each domain cluster and converge. That is, although both natural kind and artifact concepts can be represented as networks of dependencies, it might be that for natural kinds these dependencies point inwards (e.g. animals have properties to serve their internal purposes), while for artifacts they point outwards (e.g. artifacts have properties to serve an external agent). Such clustering may reflect the belief that natural kinds have

essences or vital forces, while artifacts do not (e.g. Atran, 1990). Figure 7.1 above represents such a state of affairs (loosely adapted from Keil, 1995, p.237).

Keil (1995, p.250) presents two types of developmental evidence in support of that distinction: (1) that even 3-year-olds are sensitive to causal direction - they understand the difference between having a property for oneself (e.g. a rose has thorns to protect itself) rather than for someone else (e.g. barbed wire has thorns to protect us), and (2) that even for preschool children it makes sense to ask what an artifact is for (e.g. what is a hammer for?), but not what a natural kind is for (e.g. what is an oak tree for?). Such 'domain' differences leave our hypothesis unaffected¹¹. They just suggest different sorts of dependency clustering in different category domains. Given that for natural kinds the dependencies converge inwards, then molecular features (e.g. for carrot: it is made out of cellulose) should be more conceptually central. Given that for artifacts they converge outwards, functional features (e.g. for bicycle: it has a seat) should be more conceptually central.

¹¹ Even such 'domain' differences should be, at best, very general. As stated in the introduction, the features of some natural kinds (like designer plants) seem to require reference to an external agent, and the features of some complex artifacts (like computers) seem not to require (at least not directly) such a reference.

CHAPTER VIII:
FINAL DISCUSSION

8.1 SUMMARY AND DISCUSSION OF FINDINGS

The present research had one main objective and three secondary ones. The main objective was to find empirical support for our centrality hypothesis, i.e. that the more central a feature for a concept, the higher its projectibility to other concepts that share its structure. One secondary objective was to obtain empirical support for our working assumption that the effects of conceptual centrality and variability on property induction are distinct. A further secondary objective was to examine whether conceptual centrality provides a domain-specific or a domain-general constraint. The final secondary objective was to investigate whether the centrality hypothesis operates under vague conditions. The present sets of experiments are revisited and their results are discussed with respect to each of these objectives in turn. Aiming to focus on the big picture, we have purposefully omitted some details. Such details can be found in the respective experimental chapters.

8.1.1 Objective 1: All else being equal, the more central a feature for a concept the higher its projectibility to other concepts that share its structure.

Two claims can be discerned from our centrality hypothesis the second of which depends on the first: (1) that feature centrality constrains projectibility, and (2) that this effect is mediated by the extent to which the target concept shares structure (features and dependencies) with the base concept.

To assess the first claim, we operationally defined conceptual centrality. Four operational definitions were used, the ease-of-imagining task that taps immutability, and three measures that directly derive from the Sloman et al.'s (1998) notion of conceptual centrality, i.e. that a feature is conceptually central to the extent that other (central) features depend on it. These direct measures were: centrality from a single dependency chain (Chapter III), centrality from the number of properties dependent on a feature (Chapters IV and VII), and centrality from the centrality of the properties dependent on a feature (Chapter V).

The use of multiple definitions of conceptual centrality served several purposes. For one, converging evidence for the centrality hypothesis from different operationalizations makes our claim stronger. Second, converging evidence from various *direct* operationalizations of conceptual centrality, supports the Sloman et al. notion of conceptual centrality as constraining induction. This is important in terms of future model development because the Sloman et al. notion comes with a model that implements it. Third, as will be shown shortly, results from each individual operational definition of centrality were amenable to competing interpretations. The use of multiple operational definitions helped to cancel out competing interpretations across sets of experiments. Fourth, some definitions of conceptual centrality (e.g. centrality from a single dependency chain) helped pit our hypothesis against other theories that used similar experimental designs but offered different interpretations.

To assess the second claim, we needed to operationalize the extent to which two concepts share structure. The judged similarity between the base and the target categories was taken as a surrogate of the extent to which two concepts are believed to share structure. Most of our experiments manipulated shared structure as follows: High-Similarity Same-Superordinate (e.g. lion, tiger); Low-Similarity Same-Superordinate (e.g. lion, mice); Low-Similarity Different-Superordinate (e.g. lion, falcon).

Claim 1: Conceptual centrality constrains property induction

The results of the present sets of experiments converge on the notion that conceptual centrality influences property induction. Broadly, we found that the more central a property for a concept, the higher its judged projectibility to other concepts.

Experiment 1 evidenced centrality effects on projectibility with centrality defined from the ease-of-imagining task. Experiment 2 replicated the centrality effect using the same materials. Experiment 2 also provided support for the idea that conceptual centrality as defined by Sloman et al. (1998) influences projectibility. Thenceforth, our experiments operationalized centrality directly from the Sloman et al. notion.

Experiments of Chapter III defined centrality from a single dependency chain. They examined whether centrality constrains induction above and beyond the principle of systematicity. The results showed that, as predicted, depended-on properties (e.g. causes) were more projectible than dependent properties (e.g. effects). The principle of systematicity cannot capture this finding since it does not consider the status of a property in a relation (e.g. it predicts no preference to project causes over effects). The results, however, are open to a different interpretation. Maybe the depended-on properties were more projectible than their dependent properties not because of centrality differences but because of a default assumption that the former properties are necessary but insufficient for the presence of the latter. When we stipulated, for example, that *For lions the hormone aldosterone regulates metabolism while the hormone corticosterone is regulated by metabolism*, participants may have taken the presence of aldosterone as a necessary but not sufficient condition for the presence of corticosterone. Given such an assumption, the preference to project the depended-on rather than the dependent properties is normative. Whenever the dependent property is present the depended-on property is present as well (the central property is necessary), but if the central property is deemed not sufficient, there may also be cases where the depended-on property is present while the dependent is absent.

Experiments of Chapters IV and VII aimed to rule out this interpretation by defining the central and the less-central properties in terms of the number of properties that depended upon each. Lots of an object's properties depended on the central feature while few on the less central feature. In essence, although both the central and the less central properties were depended-on, they were not parts of a single dependency chain, and hence not amenable to interpretations based around necessity and sufficiency. The results from these sets of experiments once again revealed centrality effects.

Experiments of Chapter V used a subtle definition of conceptual centrality. An equal number of properties depended upon both the central and the less-central properties (except in Experiment 12). However, the properties dependent upon the central features were more central themselves than the properties dependent upon the less central features. This

operationalization examined the iterative aspect of the Sloman et al.'s (1998) notion, i.e. that a feature's centrality is a function of the centrality of the properties that depend upon it. The results once again revealed centrality effects.

In sum, the present research strongly suggests that conceptual centrality constrains projectibility.

Claim 2: The effect of conceptual centrality on projectibility is mediated by the extent to which the base and target concepts share structure.

The findings of the current research also support the second part of our hypothesis. Critically, they suggest ways to specify our hypothesis more precisely. The "shared structure" that most seems to mediate the projectibility of a candidate feature, is the structure related to the candidate feature (such as the properties that depend upon the candidate feature.)

The series of experiments of Chapter III defined centrality in terms of a single dependency relation. In our best controlled experiments (Experiments 6 and 7) the property that depended- (or was depending-) on the candidate feature was specific and fairly general (e.g. metabolism). These experiments failed to evidence a similarity by centrality interaction. We reasoned that this may be because (1) similarity was not widely varied (all base-target pairs involved animals) or (2) the depended-on properties were fairly general and thus all target animal categories were assumed to have them. Conceptually, all target animals could be assumed to be equally similar to the base with respect to the dependent properties.

Aiming to decide, in part, between hypotheses (1) and (2), experiments of Chapter IV defined centrality from the "Few/Lots" definition. In contrast to the experiments of Chapter III, the properties that depended upon the central and the less central features were left vague (e.g. lots of a lion's functions depend upon the hormone aldosterone, but few upon the hormone corticosterone). In Chapter IV, a significant similarity by centrality interaction was evidenced. Contrary to (1), it seems that similarity was varied widely enough in Chapter III to observe such an interaction. The results of Chapter IV sit better with hypothesis (2). The effect of centrality on projectibility seems to critically depend on the extent to which the

target concept is believed to share the properties and functions that depend upon the candidate feature for the base. The results of Chapter IV can be accounted for, if one assumes that the participants used some overall similarity measure to compute the extent to which the target concept shared the properties of the candidate feature specified for the base. If they did, then the more similar the base and target concepts were, the bigger the difference of coherence lent by the central and the less central features, and thus the bigger the centrality effect, the result that was obtained.

This interpretation poses a potential threat to our claim that conceptual centrality as defined by Sloman et al. (1998) constrains inductive inference. Possibly it is not as much the centrality of the properties that depend upon a candidate inference that constrains induction, as the extent to which the target shares the features specified for the base category. The experiments of Chapter V suggested that both factors are critical. Upon both the central and the less central properties the same number of features depended, but the features that depended upon the central property were more central themselves. The results (bar the pure projectibility case, e.g. an inference from the sample of hippos to all hippos) showed main effects of both centrality and shared structure but no shared structure by centrality interaction.

In sum, the results showed that the conceptual centrality effect is mediated by the extent to which the target concept shares properties with the base and especially those properties that depend upon the candidate feature. Our findings support the following modified hypothesis:

All else being equal, the more central a feature for a concept the higher its projectibility to other concepts that share its structure, and especially the part of its structure connected to that feature.

8.1.2 Objective 2: Conceptual centrality versus variability

Our research program assumes that conceptual centrality and variability are conceptually distinct, that they refer to different perspectives of looking at concepts. We take conceptual centrality to be an aspect of the mental representations of categories, while variability to be an aspect of the set of remembered instances of a category. We expected therefore these two dimensions to be empirically distinct, to exert different forces on projectibility. To investigate this claim the experiments of Chapter V pitted centrality against variability in judgments of projectibility and judgments of frequency. The results supported our claim since they showed that: (1) the centrality of a property influences its projectibility irrespective of its objective frequency, and (2) that to a lesser extent the centrality of a property can bias the perception of frequency. Experiments of Chapter VI took this claim a step further, and asked what sort of information would people seek to make a property induction. The results suggest that centrality information has more informational value than variability information (at least for arguments with low-similar categories), since the participants preferred the former type of information over the latter.

In sum, as we expected, the effect of centrality is distinct from the effect of variability in judgments of projectibility. Interestingly, the centrality of a feature not only constrains property induction but may also bias the perceived homogeneity of a feature. Finally, centrality information is (generally) valued more highly than variability information when making an inference.

8.1.3 Objective 3: Property induction: Domain specific or domain general ?

Our account presupposes that category-based induction is influenced by conceptual centrality (a domain general aspect of representations), and therefore assumes that the centrality hypothesis will hold across different domains. We tested our hypothesis in two domains: the biological and the artifact domain. The bulk of our experiments (Chapters II-VI) used materials from the animal domain and supported our hypothesis. Chapter VII extended

the centrality hypothesis to the domain of artifacts. Taken together, the findings suggest that conceptual centrality provides a domain independent constraint on property induction.

8.1.4 Objective 4: Vagueness

The final secondary objective of our research was to examine whether the centrality hypothesis operates under conditions of vagueness. To this end, several of our experiments (Experiments 5, 8 to 11, and 17) used abstract properties or categories. Creating conditions of vagueness helped us contrast our hypothesis against structural mapping models. Notice that for structural mapping models to work, candidate properties as well as their relations to other properties need all be completely specified. That is, such models (and indeed most cognitive models) cannot operate under vagueness. However, often people make inferences under conditions of ignorance such as when we lack relevant knowledge or when we are operating in a novel context. In these cases, centrality offers a useful basis for judgments.

The centrality effect was evidenced under vague conditions. This finding gives the centrality hypothesis a heuristic status. That is, wherever possible people seem to use centrality information to make an inference. It also points to the inability of the current models of category-based inference to address centrality effects.

8.2 RELATIONS OF CURRENT FINDINGS TO MODELS OF CATEGORY-BASED INFERENCE

In the introductory chapter we presented four models of category-based inference: Sloman's (1993) feature-coverage model, Osherson, Smith, Wilkie, Lopez, and Shafir's (1990) similarity-coverage model, Smith, Shafir, and Osherson's (1993) gap model, and the structural mapping model developed by Gentner and her colleagues (e.g. Gentner, 1983, 1989). None of these models can account for the present findings.

Being similarity-based, the feature-coverage and the similarity-coverage models do not assume, or at least cannot readily capture, relational structure. The centrality effects are structural effects, and thus fall outside the scope of these models. To be fair, the centrality

effects also fall outside the domain of these models. In the Sloman and Wisniewski's (1992) sense of predicate blankness, the predicates used in the present studies were not blank since the participants could relate them back to the argument's categories.

The Smith et al. (1993) gap model also fails to capture centrality effects. Recall that the gap model deals with non-blank predicates such as *can bite through barbed wire*. Such predicates are assumed to invite an examination of the plausibility of the argument's premise. Broadly, the more implausible the premise is judged on the basis of prior beliefs, then the more these beliefs are revised as to adopt a more liberal criterion for accepting the predicate. The more liberal the criterion for accepting the predicate, the higher the judged strength of the conclusion of the argument. The predicates used here (e.g. has the enzyme streptokinase) were unfamiliar to the subjects and hence they fall outside the gap model's domain. We cannot see how the predicates used here invited an examination of the plausibility of the premises.

The structural mapping model (e.g. Gentner, 1989) can capture relational structure effects, albeit not centrality effects. To capture centrality effects a model should consider the dependency status of the candidate property in a relation - the structural mapping model does not do this. Further, the centrality effect was evidenced under conditions of vagueness, conditions under which the structural mapping model cannot operate.

8.2.1 Modeling centrality effects

At present our centrality hypothesis is just that, a hypothesis proposing a very specific constraint on property induction. It is highly desirable to develop a model to capture the present findings. There is of course the Sloman et al. (1998) model that captures the relative centrality of features for a single concept, but it is non-trivial to extend this model to category-based inference. Below we summarize the main findings of our research related to the centrality hypothesis and build up toward a simple model that can account for those findings.

Main Findings. The findings can be summarized as follows:

- (1) Experiments of Chapter III defined relative centrality from a single dependency chain (e.g. For lions, hormone aldosterone regulates metabolism, whereas hormone corticosterone gets regulated by metabolic rate). The properties that depended upon the candidate features were both specific and general (e.g. metabolism) and so all target animals could be assumed to have them. The results showed a main effect of centrality, but no shared structure by centrality interaction.
- (2) Experiments of Chapter IV defined relative centrality by the number of properties that depended upon a candidate feature. The properties that depended upon the candidate features were left vague. One can assume that the more similar a target to the base, the more likely it is that the target will share the dependent properties specified for the candidate feature. The results showed a main effect of centrality, and a shared structure by centrality interaction. The more similar the categories of an argument were judged, the higher the centrality effect; i.e. the bigger the difference between the projectibility of the central minus the less central property. The results also showed a tendency for a negative centrality effect for very dissimilar animal pairs (e.g., falcons, dolphins).
- (3) Experiments 13 and 14 of Chapter V defined centrality by the centrality of the properties that depended upon a candidate feature. The properties that depended upon the candidate feature were specific and fairly general. All target categories had them. The results showed a main effect of centrality, but no shared structure by centrality interaction.

Models. We take a property to be projectible to a target to the extent that its projection would promote coherence to the target. Adjusting the Sloman et al. (1998) notion of coherence to the present purposes, a property is projectible to a target to the extent that its projection would lend lots of structural support, that is, to the extent that it would support many of the target's (central) properties. The models we propose center around this notion.

Under the assumption that concepts are reducible to vectors of values over sets of features (as Sloman's (1993) feature-coverage model assumes), the effect of centrality might

be proportional to the extent that the target (is believed to) share the dependent properties specified for the candidate feature in the base. A simple instantiation of this idea is to take the effect of centrality on projectibility as the dot product of the vectors representing the base and target concepts with respect to the dependent properties. Call this *the similarity model of inference*.

Table 8.1. Illustration of how the similarity model can capture the shared structure by centrality interaction of Chapter IV. The number of dependent features that the target shares with the base is taken as a direct measure of the effect of centrality.

	BandT, d-on	Lots d-on	Effect on central	Few d-on	Effect on less-cen.	Central. Difference
<u>Very high similarity</u>	80%	20	16	10	8	8
<u>Low similarity</u>	60%	20	12	10	6	6
<u>Very low similarity</u>	40%	20	8	10	4	4

Note. BandT, d-on refers to the base-target feature overlap with respect to the dependent properties. We assume that this overlap will be proportional to the overall feature overlap between the base and the target concepts. Lots d-on / Few d-on stand for the number of properties that (are believed to) depend upon the central and the less central features respectively. The numbers in the example are arbitrarily chosen. The critical assumption is that the number of Lots d-on is greater than the number of Few d-on. Effect on central/less cen. refers to the effect of centrality on the central and the less-central properties, respectively. For simplicity, this effect is computed as the number of properties that the target shares with the base. Central. difference stands for the difference in projectibility between the central and the less-central properties.

The similarity model can account for findings (1) and (2). It can account for the lack of shared structure by centrality interaction in (1), because all target concepts shared the dependent properties. Table 8.1 above illustrates how the similarity model can account for the significant shared structure by centrality interaction in (2).

The similarity model, however, cannot account for the main effect of centrality in experiments 13 and 14 (finding (3)), because general features depended on both the central and the less-central properties. Hence, the amount of featural overlap with respect to these properties was the same for all target categories. Importantly, this finding shows that simple

similarity-based models will prove insufficient to account for the current findings. To capture this finding, a model needs to capture the centrality of the dependent properties. The similarity model can be extended to capture this by representing concepts as vectors over features weighted by conceptual centrality. Such feature vectors can be supplied by the Sloman et al. (1998) model. Call this *the centrality model* of category-based inference.

Such a centrality model cannot capture negative centrality effects. Such effects were observed, however, in some experiments of Chapter IV involving arguments with very low similarity category-pairs (e.g. lion-falcon). The broad shortcoming of the centrality model of inference may be that it fails to take into account the number (and type) of features that depend upon the candidate features in the base, but that the target does not have. An example would be if people were presented with a two-premise argument like the following:

Facts: Robins have property x1 upon which flying, metabolism, and heart rate depends.
 Robins have property x2 upon which metabolism and heart rate depends.

Conclusion 1: Cats have x1.
Conclusion 2: Cats have x2.

The centrality model predicts that the two conclusions would be judged as equally strong. However, it seems that the premises support conclusion 2 more than conclusion 1. The idea is: If robins have wings to fly, and flying depends on property x1, then a cat will be less likely to have x1. In general, we suggest that the projectibility of a feature might be a negative function of the dependent properties that are specified for the base but are not found in the target (i.e. **Band~T**). Using arguments similar to those in the example would help address this prediction qualitatively.

A simple way to capture such effects (of which we believe that the negative effect of centrality is but a special case) is by computing the effect of centrality on projectibility as a weighted difference between the dot product of the vectors of the base and the target categories with respect to the dependent properties (weighted for centrality) minus the dot

product of the base vector with itself with respect to the dependent properties that the base category has but the target hasn't. This effect can be computed as,

$$a (\text{BandT, d-on}) - b (\text{Band}\sim\text{T, d-on}) \quad (8.1)$$

where a, b are non-negative constants.

Table 8.2 illustrates how an instantiation of such a model can capture the negative centrality effect of Chapter IV. To simplify things, we take the overall effect of centrality in the projectibility of a feature as the difference between the number of dependent properties that the target shares with the base minus the number of dependent properties that it does not. That is, we assume that all dependent properties have equal centrality weights, and that constants a, b in Equation (8.1) equal 1.

Table 8.2 The centrality model can account for the negative centrality effect for the very-low in similarity pairs of Chapter IV.

	BandT, d-on	Band~T d-on	Lots	Effect on Central	Few	Effect on L-central	Central. Difference
<u>Very high sim</u>	80%	20%	20	12	10	6	6
<u>Low sim</u>	60%	40%	20	4	10	2	2
<u>Very low sim</u>	40%	60%	20	-4	10	-2	-2

Simple algebraic manipulations show that (under the current assumptions) negativity will be the case whenever the base-target overlap with respect to the dependent properties (i.e. BandT, d-on) is less than 50% (see Table A.8.2 for a proof).

In sum, guided by the present findings, we described a class of feature-based models that can capture the present effects of centrality on projectibility. Call this class of models the *modified centrality models* of category-based inference. At the heart of the modified centrality models lies the idea that a feature is projectible to a concept to the extent its projection would provide support to many of its (central) properties, and would leave few of its dependent

properties unaccounted for. This class of models is centrality-based because concepts are represented as vectors of values over sets of features, values that reflect centrality weights. The proposed class of models take therefore advantage of the Sloman et al. (1998) model of weighting features by conceptual centrality.

8.2.2 Toward a complete model of category-based inference

A complete model of category-based inference, to be viable, would need to also account for other known effects of category-based inference (refer to section 1.3). It would need to account, for instance, for the observations that the premise-conclusion similarity, the premise typicality, and the conclusion-homogeneity all constrain induction. Modified centrality models can capture premise-conclusion similarity effects since one of their variables is premise-conclusion featural overlap. These models, though, fail to capture typicality and conclusion homogeneity effects. In fact, for positive values of b it seems that they get typicality effects the wrong way around (See Equation 8.1.). They seem to predict, for instance, that inferences from typical to atypical categories (e.g. an inference from falcons to bats) are weaker than inferences from atypical to typical categories (e.g. an inference from bats to falcons). That is because bats (presumably) have more idiosyncratic properties. [It is interesting to note that the present class of models predicts argument strength as a negative function of $\text{Band} \sim T$ whereas the Sloman (1993) model predicts argument strength as a negative function of $\sim \text{Band}T$].

A simple (though not elegant) way to capture both category effects and centrality effects on projectibility is by linearly combining the modified centrality models of inference with Sloman's (1993) feature-coverage model. To keep it in line with the present representational assumptions, the input to the feature-coverage model would also be vectors of values over sets of features weighted for centrality. When no centrality information is provided, the hybrid model would reduce to the feature-coverage model. Since the feature-coverage model can capture the 3 main effects of categories on inference (see section 1.4.1), so will such a hybrid model. Such feature-coverage centrality hybrid models are promising in

that they seem able to account for many known effects of category-based induction. At the same time they are uninteresting because neither of the component models is an elaboration of the other.

What about systematicity effects? The class of modified centrality models, being based on the Sloman et al. (1998) centrality model, fail to capture the effect that related features are more projectible than unrelated features (but see Lassaline, 1996; Wu & Gentner, 1998) - they just consider the dependent properties as relevant. Capturing such effects is even more pressing because experiments 6 and 7 in Chapter III showed a very small difference between the projectibility of the depended-on and the dependent properties. A way to capture such effects might be by modifying the Sloman et al. (1998) centrality model by allowing some activation to flow back from the depended-on to the dependent properties. By doing so, the resulting values of the 'effect' features would be higher than the values for unrelated features. Such a breed of centrality models of inductive inference would show a preference to project related over unrelated properties.

Alternatively, one could start with a structural mapping model and modify it by adding a centrality constraint - a constraint that will bias the model towards projecting central properties. Such a modified structural mapping model would readily account for the finding that the extent to which the target shares the dependent properties of the base mediates the effect of centrality. That is because the more dependent properties the target shares with the base the better the alignment, and hence the stronger the sanction to carry over the candidate feature to the target. Such a modified structural mapping model would also be able to capture some of the main effects of categories on category-based inference (see section 1.4.2). However, it is not obvious how structural mapping models can be modified. Further, it is doubtful that such a resulting model would work under vagueness. A decisive determinant, perhaps, between modified feature-based or modified structural-mapping models would be whether the labels of relations matter (beyond merely influencing the strength of the dependencies). If they do, then the structural mapping models would be favored. If they do not, then the simpler feature-based models can capture the effects.

To conclude, we have shown that modified centrality models of inference can capture our centrality findings. Further, we have argued that hybrids combining such models with Sloman's (1993) feature-coverage model, may account for many of the known effects of category-based inference. Such hybrids are uninteresting since the component models do not build on one another. It is the same as saying that when centrality information is available, the centrality model will take over and for blank predicates the feature-coverage model will take over. Modifying structural mapping models so as to capture centrality effects offer an alternative approach toward a universal category-based model.

8.2.3 Relations between category-based models

Domain relations. What is the relation underlying the various models of category-based inference? As Gentner and Medina (1998) propose, the relation may be this: in the absence of sufficient knowledge (such as in the case of blank predicates), people fall back on default strategies such as overall similarity measures and or diversity-based reasoning. In such cases, both the similarity-coverage and the feature-coverage models are adequate for capturing judgments of projectibility. When knowledge is present relating the candidate feature to other properties of a concept, people will make intensive use of such knowledge. In such cases judgments of projectibility will be constrained by structural aspects such as the principle of systematicity and the centrality status of the candidate feature (For corroborating evidence that people make intense use of represented knowledge whenever possible, see section 8.6.4 on Instability across cultures.). In such cases, the structural-mapping and the centrality-based inference models will be applicable. In the special case where such knowledge is vague, a vague centrality heuristic would be useful. Such cases are in fact not that special. We often make inferences under conditions of ignorance such as when we lack relevant knowledge or when we are operating in a novel context (see also section 8.6.3 Instability across development). Centrality models of property induction could best capture such cases. Finally, in cases where the plausibility of premises is at question, the gap model may prove the best. In other words, the centrality heuristic is but one strategy that people use

to infer properties. That is why, we believe, it is uninteresting to try to combine centrality with models about blank predicates.

Maximizing coherence. As stated in section 1.4.3, a central principle underlying category-based induction (and concept use in general) might be that of increasing consistency (positive coherence) and decreasing or avoiding inconsistency (negative coherence). (For more on positive and negative coherence see Pollock, 1979.) Property generalization is by itself a means of increasing positive coherence. Projecting a property from a sample of instances to the whole category can be seen as an inference to the best explanation, taking the generalization to explain its instances (see Harman, 1995). This explains, in part, why we have a tendency to draw inferences in the first place. One of the reasons why we endorse some property generalizations more than others seems to be because some properties lend more coherence to the target concept. That is, the purpose of property induction might be to increase the target concept's coherence. Projecting properties that are thought to explain (or to provide structural support for) many of the target concept's properties promotes more coherence than projecting properties that explain fewer of the target's properties. Projecting based on matching systems of related rather than of unrelated properties (i.e., following the principle of systematicity) can be also seen as a direct way of maximizing the target's coherence (see Harman, 1995).

Stating the problem of induction in terms of maximizing coherence leads us to expect certain other aspects to potentially constrain induction, those aspects that have been proposed to influence coherence. Harman (1995) proposes that a coherence-giving connection orthogonal to explanatory connections, is implication. In the abstract, say that you believe that either concept C has feature A or else that it has feature B. Your senses inform you that concept C does not have feature B. Then the implication that C has feature A increases the concept's coherence. Harman suggests that implication cannot be reduced to explanatory relations (to say that in the above example your beliefs and observations explain the conclusion seems to stretch the notion of explanation). Implications, though, might be

reducible to asymmetric dependency relations, and hence they might be captured by the Sloman et al. (1998) notion of conceptual centrality.

SUPERORDINATE STRUCTURE

8.3.1 Hierarchies, judged similarity and property induction

As noted in the introductory chapter, several theorists assume, either explicitly or implicitly, that (at least animal) concepts are structured hierarchically (see e.g. Osherson et al's (1990) similarity-coverage model). As a result, such theorists claim that people should be more willing to endorse inferences among categories from the same close superordinate than from a different superordinate (for empirical evidence see Gelman & Markman, 1986; Gelman, 1988). An inference from a mammal to another mammal, for example, should be judged stronger than an inference from a mammal to a bird. We argue that this may not always be the case. First of all, this assumption may fail with nonblank candidate predicates. Predicates such as *can bite through barbed wire* have been shown to potentiate some of the premise category's features. The inductive strength for such nonblank properties has been found proportional to the extent that the target category is similar to the base with respect to the potentiated properties (see the Smith et al. (1993) gap model). Conceptually, one can think of many cases where two categories from the same superordinate do not share the critical properties, whereas two categories from a different superordinate do. In general, as Goodman (1972) pointed out any two entities may be maximally or minimally similar depending on the chosen "respect". Consider the following categorical arguments:

- A. Cats can fit in box A; therefore pennies can fit in box A.
- B. Cats can fit in box A; therefore tigers can fit in box A.

Argument A seems much stronger than argument B though the former involves an inference among categories from a different superordinate whereas the latter an inference among categories from a common superordinate.

More interestingly, even unfamiliar predicates seem to potentiate some of the base category's properties. Heit and Rubinstein (1992) showed, for instance, that for the anatomical property *has a liver with two chambers* inferences were stronger from hawks to chickens than from tigers to chickens, whereas for the behavioral property *prefers to feed at night* the order of the preference was reversed. Hence, the claim that sharing a superordinate plays a special role in property induction seems unwarranted.

In a similar line, Experiments 6, 7, 10 and 11 failed to reveal a special status of superordinate structure on inductive inference. No robust difference was found in the projectibility of properties among LSSS and LSDS categories. At the same time, a significant difference was found in the projectibility of properties among HSSS and LSSS and among HSSS and LSDS categories. The present experiments suggest therefore that judged similarity between the premise and conclusion categories rather than superordinate structure decides the projectibility of a property.

8.3.2 Hierarchies, judged similarity and the effect of centrality on projectibility

Relating hierarchical structure to the centrality effect, one might suggest that the centrality effect is proportional to the extent that the base-target categories share a common superordinate. Conceptually, this should not always be the case at least for nonblank predicates for similar reasons as those outlined above. Empirically, the results of Experiments 6, 7, 10 and 11 speak against such a claim. Experiments 6 and 7 showed a significant main effect of centrality but no shared structure by centrality interaction. Sharing therefore a superordinate failed to decide the absence or presence of a difference in projectibility between the central and the less central properties (i.e. the presence or absence of a centrality effect). Experiments 10 and 11 showed a main effect of centrality and a significant shared structure by centrality interaction. Specifically, the effect of centrality on projectibility was

proportional to the judged similarity of the base-target categories. The present experiments therefore also rule out the possibility that superordinate structure decides the presence or absence of a centrality effect on property induction.

8.3.3 Hierarchies, dependency structure, centrality and property induction

A more sober claim might be that categories inherit dependency structure from their superordinates. If this were true, then a given property present in two categories from the same superordinate should have a comparable centrality status in each. We claim, however, that dependency structure may differ within a hierarchy and because of that a given property may have different centrality statuses across categories from the same superordinate. Take, for instance, wings. Both falcons and chickens have them, but for falcons wings are more central property since flying (and all the properties that depend on flying) depends on them. In essence, as Sloman et al. (1998) proposed, we hold that the centrality of a property directly derives from the number and the centrality status of the features that depend on it. Since more features about falcons than about chickens depend on having wings, having wings is a more central feature for the former than for the latter category. In addition, given that wings are more central for falcons than that for chickens, the present hypothesis predicts that wings should be more generalizable to a new species of falcons than to a new species of chickens.

We put our thought experiment to the test. Thirty-seven participants were asked to rate the centrality of having wings for falcons and chicken. To the extent that the Sloman et al. (1998) model is right, we expected participants to judge wings to be more conceptually central for falcons. The same participants were also asked to estimate the likelihood of encountering a wingless species of falcons (chickens). To the extent that wings are more central for falcons and that our centrality hypothesis is right, we expected people to rate the possibility of wingless chickens as more likely. The results supported both claims (see Table A.8.1). Wings were judged to be marginally more central for falcons than for chickens ($p < .06$). Furthermore, participants rated as more probable to encounter a species of wingless chicken than a species of wingless falcons. This preliminary study showed therefore that (1)

dependency structure may differ within a hierarchy, (2) the centrality status of a property is a direct consequence of the dependency structure in which it is embedded, and (3) that centrality influences property induction. Last but not least, the present study suggests that our centrality hypothesis extends to familiar properties and categories. This last point is further reinforced from the Experiments of Chapter V, especially Experiment 14, which showed that familiar central properties (e.g. heart, a strong immune system) are more generalizable than familiar less central properties (e.g. mud spots).

8.4 SAMPLE SIZE AND MULTIPLE PREMISE ARGUMENTS

The present research is limited in at least two respects: it does not address effects of sample size or effects of multiple premises. Below we present some novel empirically testable predictions for each case.

8.4.1 Sample size

Nisbett, Krantz, Jepson, and Kunda (1983) presented evidence that inductive inference is sensitive to sample size to the extent that the candidate property is *believed* to be somewhat variable. Participants were informed, for instance, that a sample of individuals of a tribe was observed all of who were obese and had brown skin. The size of the observed instances (1 or 3 or 10 individuals) was manipulated between participants. Nisbett et al. found that the effect of sample size on projectibility was high for properties believed to be heterogeneous (e.g. obese) but not for properties believed to be homogeneous (e.g. brown skin).

Nisbett et al.'s research did not address why, in the first place, some properties are believed to be more homogeneous than others. Specifically, it may be that these notions arise bottom-up from frequency tabulations (this idea cannot apply to their examples since the context was novel), or top-down from expectations based on prior knowledge. Our research presents evidence that the centrality of a property may influence, in a top-down fashion, its

projectibility; in fact we have shown that centrality can bias perceived homogeneity (see Chapter V). Recall that most of our experiments used novel properties thus precluding importing frequency tabulations from similar categories. To the extent that central properties are (all else being equal) more projectible, and given that the more projectible a property the less the effect of sample size, we predict that the projectibility of central features would be more robust with respect to sample size.

Here is what we have in mind. Even if we just find out that a single PC has memory registers upon which the functions of its central processor depend, we expect people to be highly confident that other PCs will also have memory registers. As more positive evidence comes in, the initial level of confidence should not increase by much. That is, even with few cases the level of confidence should be close to ceiling. If we find out, however, that a single PC has a red cable that almost none of its features depend on, we will be wary of inferring that other PCs have red cables as well. As more positive evidence comes in, say 20 out of 20 PCs have the red cable, we will be much more willing to generalize this property to other PCs. As Keil et al. (1998) suggest, in the absence of theories, feature tabulations serve as data for theory development. Returning to the example, people might even try to explain the apparent feature homogeneity. They may reason, for instance, that few companies make this sort of cable and they all make them red, or that there might be a reason for why the cable is red that we have not yet discovered.

8.4.2 Multiple Premise Arguments

At present our studies are also limited to single-premise categorical arguments. It is reasonable to wonder how our theory extends to multiple-premise arguments. On a similar line of reasoning as with sample size, we believe that (all else being equal) the number of premises will exert a stronger influence on the projectibility of less-central properties. An example will help illustrate. Given that lions have a property upon which some of its central functions (such as the amount of oxygen in the blood and metabolic rate) depend, we expect participants to be highly confident that tigers will have the property as well. Stating that other

categories (such as jaguars, pumas, and hyenas) also have the candidate property, should not increase their confidence by much. Their confidence should be close to ceiling even with the single premise argument. In contrast, the fact that lions have a less central feature (e.g. a hormone that controls teeth shine) does not strongly sanction that tigers will have the property as well. Stating that other categories have the property as well should substantially increase belief in the argument's conclusion¹².

8.5 IS FOLLOWING THE CENTRALITY HEURISTIC ADAPTIVE?

The evidence that centrality constrains projectibility even under conditions of vagueness, gives the centrality hypothesis a heuristic status (see the section on 8.1.4 on Vagueness). A critical question is whether following the centrality heuristic (i.e. whether projecting conceptually central over conceptually less central properties among concepts to the extent that they share structure) is an adaptive strategy.

Adaptiveness is a pluralistic term - being adaptive means different things to different people. Whether projecting properties according to our hypothesis is adaptive depends on the particular sense of adaptiveness adopted. Below we unpack two senses of the word "adaptiveness" and evaluate whether following the centrality strategy is adaptive in each case. But first we argue that our working assumption that concepts are organized around central features is adaptive in the sense that it help us master our environment.

The idea that central properties are weighted more heavily in inductive decisions is based on the notion that concepts are structured around central features. Is it adaptive to organize concepts around central features (e.g. causes rather than effects)? We argue that it is adaptive because, as Ahn (1998) proposes, it increases our chances of survival; it helps us master our environment by promoting predictability and controllability. Its adaptive value

¹² Some researchers do not differentiate between generalization from instances and category-based inference; they treat both as generalizations from a set of cases (e.g. Heit, 2000). A critical difference between the two might be that instances of a category are more homogeneous than different categories. This implies that the projectibility of central properties is more robust to the number of "cases" when generalizing from instances than when generalizing from categories.

also seems to be domain independent. It makes sense to class together biological organisms that share conceptually central rather than conceptually less central features. Animals sharing conceptually central features (e.g. being a carnivore) helps us infer that they also share behaviors (e.g. being a predator, being relatively agile) as well as other biological features (e.g. presence of some enzymes, high rate of digestion, big field of stereoscopic vision). Similarly, it makes more sense to group disease concepts around etiology, rather than symptomatology. Classifying diseases according to etiology (but not symptomatology) helps predict the progress of the disease and to prescribe a treatment. Social concepts are also more inductively potent if grouped in terms of shared central features (such as psychological traits) rather than shared non-central features (such as behaviors). Psychological traits (e.g. being hot tempered and aggressive) help us predict a person's behavior in a novel situation (e.g. likely to initiate a fight in a bar), whereas behaviors (e.g. fighting in bars) do not readily generalize to novel situations, such as behavior at work.

In sum, it makes good sense to organize concepts around central features because doing so enhances their inductive power. Below we argue directly that projecting central over less central features across similar concepts is an adaptive strategy in terms of both maximizing real world success and minimizing psychological conflict.

8.5.1 Adaptiveness as maximizing real world success

One way to measure the extent to which a strategy is adaptive is by looking at its relative success rate. Would one be most times right in projecting central over less central properties to the extent that the base and the target categories are similar? First off, recall that centrality is not only concept relative but also relative to an individual. A feature's conceptual centrality is conditioned on a person's knowledge and intuitions about how the various features of a concept depend on one another. If part of this knowledge is fallacious, then properties that are assumed to be central will not be. If one believes, for instance, that lots of properties about refrigerators depend on *being white* but few on their ability to *freeze up stuff*, then for that person color would be a more immutable feature of refrigerators than function.

Most people, however, have a good idea of how properties are intertwined (at least for objects that we are relatively familiar with); this is why the belief that whiteness is a relative central property of refrigerators sounds strange. Given that most peoples' theories are correct, the degree of mental transformability should reflect the degree of actual transformability across members of a category; properties that are thought central should be more homogeneous. For example, it is more likely to find a pink refrigerator than one that was not designed to freeze up stuff. Similarly, one should be more confident that members of a similar category, such as freezers, will freeze up stuff rather than they will be white.

The above argument is grounded on intuition. Corroborating evidence that following the centrality strategy is adaptive, at least for the biological domain, comes from the famous biologist and statistician Ronald Fisher. By means of an analogy, Fisher argued that evolution progresses gradually rather than in big steps. At a given time in evolution a species can be seen as analogous to a microscope slightly out of focus. This analogy goes through because a species, to have survived so far, should be close to perfection, "slightly out of focus". Fisher claimed that further improvement could be best achieved by relatively small changes (micro-mutations) rather than by relatively big changes (macro-mutations). Relative small changes have about 50% chance of being successful. Big changes have a lower probability of success. To better understand Fisher's claim assume that the microscope's initial state is 1 inch out of focus. Making less than 1 inch changes in either direction has a 50% chance of being successful, of bringing the object in focus. Making changes bigger than 1 inch has a much less chance of success. This illustration presupposes knowledge of the magnitude of the initial error. The beauty of Fisher's argument is that it stands independent of this knowledge. Even when the amount of initial error is unknown, the smaller a change the closer the chances of success to 50%, the bigger the change the further away the chances of success from 50%. (For a thorough exposition of Fisher's argument see Dawkins, 1986).

In our terminology, changing a central property constitutes a macro-mutation whereas changing a less central property a micro-mutation. That is because changing a central property forces a cascade of other feature mutations. For the mutated species to be successful,

these other mutations should be in the desired direction as well. The greater the number of successful mutations required, the less the likelihood of overall success¹³. Returning to our main issue, Fisher's argument translates that into the idea that the closer two species are together in evolutionary terms, the more central properties they share. Hence, all else being equal, projecting a central over a less central property to the extent that the base and the target animals are similar is an adaptive strategy since it correctly captures the structure of the environment.

Applying Fisher's argument to our purposes presupposes that judged biological similarity captures evolutionary closeness. Is there any evidence that people's judgments of similarity reflect to some extent closeness in evolutionary terms? A strong case can be made from development and cross-cultural evidence. [Evidence from adults of Western cultures would be weak since it may reflect the impact of education]. We begin with the developmental evidence first.

Preschoolers base their inductions on category membership even when this goes against perceptual similarity (e.g. Gelman, Collman, & Maccoby, 1986; Gelman & Markman, 1986, 1987). Gelman and Markman, for instance, presented 3- and 4-year olds triads of pictures, 2 base pictures and 1 target picture, each featuring an object. Children were taught a novel fact about each of the base objects and were asked to generalize one of these facts to the target object. The stimuli were manipulated such as the target object looked very much like one base object but shared the same category label as the other base object. For example, children were presented with a tropical fish and were told that it breathes underwater. They were also presented with a dolphin and were told that it pops out on the surface to breathe. They were then asked to decide which of these 2 facts generalized to a shark (which looks more like a dolphin but belongs to the same category with the tropical fish). The results showed that children's inductions frequently relied on category membership rather than

¹³ This argument can be made even stronger if we take into consideration that more mutations require more effort. The expected value (expected effort times expected probability of success) therefore of macro-mutations is much less than the expected value of a micro-mutations.

perceptual similarity. The use of common labels cannot account for the results. Even when no labels were used, children often drew inferences based on category membership (Gelman & Markman, 1987). So, developmental evidence suggests that preschoolers' inductions are influenced to some extent by evolutionary closeness.

The preceding developmental evidence comes from a select group: American children. As Super (1980) argues, cross-cultural research can inform developmental theories by expanding our notions of what constitutes the endpoint of development. Basing theories on a select group of people, runs the risk of underestimating the variability of the adult end states, and thereby oversimplifying accounts of conceptual development. Recently, cognitive scientists have begun to study reasoning and categorization across cultures (e.g. Atran, 1995; Atran, Estin, Coley, & Medin, 1997; Coley, Medin & Atran, 1997; Lopez, Atran, Coley, Medin, & Smith, 1997). Lopez et al. (1997), for instance, studied the organization of folk biological taxonomies and their use in inductive reasoning among American and Itzaj Mayan adults. They found that the folk biological taxonomies of these populations were very closely and to a similar extent related to scientific categories. Lopez et al.'s experiments also suggested that a universal feature of inductive reasoning is making use of a culture's folk taxonomy. The fact that the folk taxonomies of these very different populations were highly related to scientific taxonomies (which capture evolutionary closeness), coupled with the fact that individuals of each culture depended on their folk biological taxonomy to infer novel properties, strongly suggests that reasoning about biological categories is influenced by evolutionary closeness.

Returning to our main argument, following the centrality heuristic for biological categories seems adaptive because it seems to capture the structure of the environment. That is because, (i) animals close in the evolutionary chain *do* share lots of their central properties, and (ii) people's folk biological inductions are sensitive to evolutionary closeness.

8.5.2 Adaptiveness as minimizing psychological conflict.

Drawing an inductive inference involves updating one's beliefs - induction is a special type of belief revision. Quine (1961) noted that logically speaking all properties of a concept are revisable. According to his view, it is no less illogical to stick to central rather than to less central properties. However, Quine also noted that features are not equally revisable psychologically; people are less willing to give up properties upon which lots of their beliefs and ideas depend. All else being equal, sticking to beliefs one cherishes is an adaptive way to update one's beliefs - being conservative minimizes changes in one's knowledge base. An example will help illustrate. Say that you cherish your belief in science; lots of your ideas, beliefs, attitudes, and behaviors center around this belief. Assume, for instance, that you spend hours reading about science and you have even written books about it. You walk in the street and you see a magician. He performs a trick that you cannot readily explain. This will (presumably) induce a negative state because your beliefs and behaviors about science are incongruent with your current experience (Festinger, 1957; see also Cooper & Fazio, 1987). What beliefs should you revise? Should you give up your beliefs and attitudes toward science and confer that you have just observed a miracle? It makes more sense to stick to your cherished beliefs and conclude that the magician performed a trick that you cannot readily explain. Doing so achieves cognitive consistency at a minimal psychological cost. (For a detailed account of why conservatism is more reasonable, in the ordinary use of the term reasonable, than competing accounts such as Descartes' (1637) "special foundationalism", see Harman 1995).

Returning to our main argument, following the centrality strategy seems to be *psychologically* adaptive. Induction is a special case of belief updating where new knowledge gets generated. In induction the question is not which beliefs should get ditched in the case of a conflict, but rather which new beliefs should get added to our existing stock of knowledge. Our knowledge about many objects is incomplete. Most lay people (ourselves included) know, for instance, that some biological substances regulate the digestion of lions though we

may not know which. In such cases it can be said that we have some *essence placeholders* waiting to be filled-in with a theory (Medin & Ortony, 1989). We argue that such placeholders are more readily filled-in by projecting central properties from similar concepts. That is because such properties would explain (or support) more of our other beliefs about the target concept.

In sum, conceptual arguments show that following the centrality strategy is adaptive because it helps us maximize both real world success and psychological coherence.

8.6 ON CONCEPTUAL ASSUMPTIONS

Our centrality hypothesis is grounded on the presupposition that concepts (the 'ideas' describing categories) involve stable mental representations. Given that such representations are both stable and reducible to sets of interdependent features, a centrality measure can be defined over them. On the one hand, the assumption of concept stability is implicit in most accounts of concept structure and use - few theorists have questioned its soundness (for notable exceptions see Barsalou, 1983, 1987, 1989; James, 1950; Smith & Samuelson, 1997). On the other hand, various instabilities concerning conceptual judgments have been reported in the literature such as judgmental instabilities across points in development, cultures, tasks, and so on. The next section explores whether such judgmental instabilities imply unstable mental representations. In a sense, the next section is about the minimal conceptual assumptions needed for our hypothesis to go through, and about whether the evidence supports these minimal assumptions. But first we discuss a particular source of instability known as the multiple descriptions problem.

8.6.1 The multiple descriptions problem: Given that an object can be described in multiple ways, does it make sense to talk about the centrality of a feature for a given concept?

Entities do not map in a neat one-to-one fashion onto categories: A given entity may belong to many different categories (to the extensions of many different concepts). At

different times, for instance, an artifact can be described as an ashtray, a smoker's paraphernalia, a paperweight, or a piece of glass. Similarly, a natural kind can be described as a dog, a German shepherd, a pet, or a friend. Some of the categories that a given object belongs to are interrelated in a special way: one category is a subset of the other - German shepherds are kinds of dogs, and ashtrays are kinds of smokers' paraphernalia (for a review on the hierarchical structure of concepts see Murphy & Lassaline, 1997). A feature's centrality status across such concepts (e.g. dog and animal) seems to be highly related: features central for the concept animal seem also central for the concept dog. Hence, at least from the outside, the fact that objects can be described at various levels of abstraction does not seem to threaten the centrality notion (but see section 8.3.3).

An object though may also claim joint membership in unrelated categories (e.g. it may be a member of the category glass and the category paperweight). Features that are central for the concept glass (e.g. being made out of glass) are not central for the concept paperweight. This type of description variability seems to threaten most the centrality hypothesis. The problem is this: given that an object can be described in many different ways and given that some of these descriptions imply very different features and dependencies, then the notion of a feature being central for an object seems unattainable. In fact such an assertion is fallacious, but this is not what we claim.

Our centrality notion is about *concepts* ('ideas' describing categories) and not about objects. The centrality of a feature for an object is therefore conditional on the perspective one adopts about that object - the mental representation that one accesses upon perceiving that object. In a context where an object is thought of as a paperweight, features such as weight (and probably size) will be judged as conceptually central. In another context where the same object is thought of as an ashtray, features like *being able to hold cigarettes* and *being made out of a non-burning material* will be conceptually central. The notion of conceptual centrality circumvents (or better, is independent of) the problem of the one-to-many mappings from objects to categories. Our centrality notion is conditional on a person's conception of an

object and not on the object itself. It makes sense therefore to talk about the centrality of a feature for a concept (as we do) but not about the centrality of a feature for an object.

Some theorists take the idea of the one-to-many mapping to an extreme and argue that there are infinite ways of grouping objects into categories - this view is known as *promiscuous realism* (Dupre, 1981). Although such a view does not contradict our thesis, it seems to question whether there are principled ways of organizing entities into categories. Keil (1995) suggests that from the premise that an object may belong to many categories it does not follow that such groupings may proliferate endlessly (at least for natural kinds). We also believe that the mapping from objects to categories is relatively constrained.

Supportive evidence comes from the fact that it is very difficult to think of artifacts in multiple ways; the conventional use of artifacts greatly constrains our ability to perceive them as suitable for a wide variety of purposes. Pliers, for example, are conventionally used to cut stuff. Although they can certainly be used as a weight, it is hard to perceive them as such. In fact, knowledge of the conventional function of artifacts has been shown to interfere with problem solving. This phenomenon is known as *functional fixedness* and was the subject of many studies (e.g. Adamson, 1951; Birch & Rabinowitz, 1951; Duncker, 1945; Maier, 1931). In Maier's (1931) classic study, participants had to tie together two strings that were hanging from the ceiling. The strings were placed far apart, so a participant could not grasp both strings at once. A participant could use a variety of objects to solve the task, including a chair and a pair of pliers. The solution involves tying the pliers on one string, setting that string swinging like a pendulum, and holding the second in the center of the room until the other string swings close enough to grasp. Only 39% of Maier's participants solved the problem within 10 min, presumably because of the difficulty of perceiving the pliers as weight.

More direct evidence for functional fixedness comes from Duncker (1945; see also Adamson, 1951). His 'box' problem asked participants to mount three candles vertically on a screen, at a height of about 5 feet. To accomplish the task, participants were given a large number of objects including three pasteboard boxes of varying sizes, five matches, and five thumbtacks, the crucial objects for solving the task. The solution was to mount one candle on

each box by melting wax on the box and sticking a candle to it, and subsequently to tack the boxes to the screen. Most people found it hard to use the box as a platform for the candle, presumably because boxes are conventionally used as a container. The functional fixedness was more pronounced when the box was initially full rather than empty, presumably because being full reinforces its function as a container.

In sum, we argue that our hypothesis is not suspect to a proliferation problem because it can be conditioned on the perspective that one adopts about an object. Further, it is not totally useless to state that some properties are more central than others for an object, because most of the times and for most people the perspective one adopts about an object is limited. It makes sense, for instance, to claim that the ability to freeze up stuff is more central for a refrigerator than being white since in the conventional perception of refrigerators the former property is more central than the latter.

8.6.2 Instability across tasks: If the relative centrality of a feature is task dependent, does it make sense to talk about feature centrality?

A robust finding in the concept literature is that the judged importance of a feature for a concept is largely conditional on the task used to measure it (e.g. Barsalou, 1983, 1989; Rips, 1989; Malt, 1994; Sloman et al., 1998). Prima facie such findings seem to challenge the notion that a single stable representation is accessed in all these tasks. We argue that such variability in judgments might instead be because different types of knowledge are important in different types of tasks (cf. Sloman et al., 1998). But first we begin with empirical evidence supporting the instability of a feature's importance across tasks.

Malt (1994) gathered judgments about the appropriateness of the label "water" for various types of liquids, the percentage that each liquid was believed to contain H₂O (presumably water's essential element), and the typicality of each liquid with respect to the category water. Malt's major finding was that category judgments were task dependent (see also Rips & Collins, 1993). Mineral water, for instance, was called water, was judged to have a high percentage of H₂O, but was also judged to be an atypical instance of the category

water. In contrast, ocean water was uniformly called water, was judged to be highly typical of water, but was judged not to contain high amounts of H₂O. Further, tea was judged to contain a higher percentage of H₂O than ocean water, but it was not even called water. At an extreme, there seem to be as many different patterns of category judgments as there are tasks to measure them. (Below we return and evaluate this extreme claim).

Sloman et al. (1998) demonstrated a dissociation between judgments of naming and judgments of mental transformability. Being "red" is important for naming something a red apple (e.g. it is highly inappropriate to call a green apple "red") but not for imagining it (e.g. it is relatively easy to transform a red apple into one that is green). Sloman et al. attributed their findings to the view that different aspects of conceptual structure are important for naming and for conceptual transformation. They suggested that the importance of a feature critically depends on the *goal* of the agent using the concept. A goal of naming is to distinguish a category from similar categories. Color, for instance, is a relatively important feature for naming apples since it helps differentiate between types of apples. In contrast, property induction (presumably) aims in increasing the coherence of the target concept. Color is not a conceptually central property of apples (not many apple properties depend on it) and hence not highly projectible. In fact, for highly similar concepts (e.g. tigers and lions) features high in immutability (e.g. heart, liver) have low discriminatory value.

In sum, the fact that the importance of a feature is task-dependent is not inconsistent with the claim that a single mental representation underlies various conceptual judgments. As Medin and Ortony (1989, p. 191) suggest, one should not equate instability in outputs with internal instability. It might be that mental representations are in fact stable and that the apparent instability is an artifact of the processes that operate on stable concepts. On similar lines Komatsu (1992, p.501) remarks that "experimental results do not directly demonstrate anything about conceptual representations. It is always difficult to decide whether a particular observation is a function of the information represented by the concept, the structure, or the format of that information". He concludes that what such results at a minimum specify is what sort of information is available to people about categories.

Parenthetically, we would like to mention that types of feature importance do not proliferate endlessly. Evidence suggests that different types of tasks load systematically on different factors, on different aspects of concept structure. Sloman et al. (1998), for instance, gathered conceptual judgments from 10 different tasks and found that the resulting judgments loaded on just 3 factors: feature centrality, feature diagnosticity, and feature salience (for more information see Sloman et al. 1998, Study 1).

8.6.3 Instability across development

The mental representations of categories change with time. Say, for instance, that the first 5 Frenchmen I meet are all rude. Based on this information, I am led to expect that the next Frenchman will also be rude. If the following 15 Frenchmen I see are in fact polite, I may give high chances that the twenty-first Frenchman will be polite. The question of whether concepts change with time seems therefore silly (for a review of the conceptual learning literature see Heit, 1997). Children's concepts are certainly different from adults', often in potentially uninteresting ways.

A more interesting question to ponder is whether children's and adults' concepts differ qualitatively. Early developmental research seemed to support such a distinction; it supported a shift from similarity to rules or from association to definition. One type of evidence came from studies showing that concepts follow a characteristic to a defining shift (Keil & Batterman, 1984; Keil, 1989). It has been found that young children judge category membership on the basis of typical or characteristic features even in the absence of defining features. Young children, for instance, judge a person of about the same age as their father, who gives them gifts as an "uncle", but not an adolescent who is their father's brother. In contrast, older children and adults weight more heavily defining rather than characteristic features (Keil, 1989).

As Keil, Smith, Simons, and Levin (1998) suggest, on a closer look, such evidence does not really support a similarity to theory shift. Even young children's concepts must be constrained by theoretical (or pre-theoretical) principles. A young child, for instance, would

not consider features such as "having blue shoes" as important for claiming membership in the category "uncle" even if all her uncles happened to have blue shoes (for strong arguments against a developmental similarity to theory shift see Keil et al., 1998). We concur with Keil et al. that even young children's concepts are constrained by intuitive theories (or at least by vague pre-theoretical biases) that some features are more explanatory relevant than others. Their arguments are compelling. Intuitive theories or pre-theoretical biases, for instance, seem necessary for constraining the otherwise vast tabulation space.

Interestingly, the sort of pre-theoretical biases (or notions of causal powers) that Keil et al. (1998; see also Keil, 1995) advocate point to a potential advantage of Sloman et al.'s (1998) notion of conceptual centrality over Ahn and Lassaline's (1995) causal status hypothesis: pre-theoretical biases can be readily captured by asymmetric dependency relations but not by causal (or generally by explanatory) relations. Explanatory relations seem to presuppose knowledge of causal mechanisms. To say that feature A explains feature B implies that one has an intuitive theory stating why this is so. Pre-theoretical biases are by definition vague - they refer to the *intuition* (the ungrounded belief) that some features are more explanatory and central (for a given domain) than others. Models of conceptual centrality can capture notions of causal powers while the causal status hypothesis cannot capture other dependency relations. Given that people's concepts are influenced by pre-theoretical biases, the conceptual centrality hypothesis would therefore have an advantage: features that by intuition seem central (and hence are psychologically central in conceptual judgments) would be represented and predicted as such¹⁴.

Evidence suggests that vague pre-theoretical biases *do* exert an early influence on judgments of category centrality and projectibility. Simons and Keil (1995) presented children with a target (an animal or a machine) together with a set of potential 'insides': the insides of an animal, a machine, a pile of rocks, or a pile of blocks. Children were asked to

¹⁴ Notice that the ability of models of conceptual centrality to capture vague pre-theoretical biases provides also an advantage over structural-mapping models. As stated elsewhere, structural mapping models cannot deal with vagueness.

match the target objects with the correct insides. The results showed that even preschool children expected the insides of machines and animals to differ (they systematically picked different insides), but sometimes they were picking the wrong insides. Vague causal knowledge seems therefore not only to influence but also to precede knowledge of specific causal mechanisms in influencing judgments of category centrality. In fact Keil et al. (1998, p.42) argue that "these basic notions of causal centrality may emerge early, as early as sensitivity to typicality of properties ... the ability to perceive and learn causal patterns may be just as fundamental as the ability to learn typicality and frequency distributions."

Further evidence directly supports the view that abstract pre-theoretical biases influence property induction. Keil and his colleagues (in preparation, cited in Keil, 1995), for instance, presented children with an ambiguous picture that for one group of children was labeled "rock" while for another "frog". Participants were shown two other pictures and had to choose the one that belonged to the same category as the target object. The results showed that when the property was labeled as "frog", shape rather than color and surface markings were important. When the same entity was labeled as "rock", color and surface markings became more important. Keil (1995) demonstrated a similar property by category interaction while using pictures vaguely described as either novel animals, or novel machines. Because such objects were by definition novel, the knowledge underlying children's judgments should be in the form of abstract pre-theoretical biases rather than of concrete causal knowledge.

To close-up the parenthesis, it seems that concepts are constrained not only by precise causal knowledge but also by vague pre-theoretical biases. To the extent that these vague pre-theoretical biases can be captured by the notion of conceptual centrality but not by the causal status hypothesis, the former has an advantage.

8.6.4 Instability across cultures: Does the observed cultural instability in concept structure and use affect the centrality hypothesis?

In the section about the 'adaptive' value of the centrality hypothesis (section 8.4), we mentioned that cross-cultural studies suggest some universals in folk-biological organization

and reasoning. For instance, Lopez, Atran, Coley, Medin, and Smith's (1997) studies with Americans and Itzaj Mayans revealed that these cultures are similar in that they both (i) organize species into folk-biological taxonomies and (ii) use their respective taxonomies to make category-based inferences. The Lopez et al. (1997) studies also revealed some cultural peculiarities. An interesting one is that while American subjects showed a tendency to judge as stronger multiple premised arguments with diverse rather than similar premises, Itzaj Mayans did not. (In the Osherson et al. (1990) terminology Americans exhibited *premise diversity* while Itzaj Mayans did not). Americans seem to have based their inductions on a vague diversity principle (something like: the more different the categories sharing a property the more likely that other categories will share the property as well), while Itzaj Mayans seem to have based their inductions on ecological considerations. Follow-up studies revealed that the lack of premise diversity in the part of Itzaj Mayans was not because Itzajs were not capable of diversity based reasoning - in different settings their reasoning exhibited diversity. Lopez et al. suggested that the premise-diversity difference might be attributable to Itzajs having more knowledge than Americans about their respective ecology. This point is supported by evidence that American experts (e.g. landscapers and maintenance workers) also show no or negative diversity as a result of engaging in ecological reasoning (Medin, Lynch, Coley, & Atran, 1997). (For people who believe that knowledge is an integral part of culture (like Lopez et al. do), such differences are cultural differences. For people that see knowledge as a factor independent of culture, such demonstrations are demonstrations of knowledge differences.)

Our centrality hypothesis is consistent with the notion that these premise-diversity reasoning differences between cultures or people of the same culture, are knowledge based. The centrality of a feature for a given person (or a given culture) directly depends on the knowledge that this person (or the culture) has. Different knowledge about a category translates into representations involving different dependency structures. The centrality of a feature across representations of different dependency structures, is potentially different. This follows from defining conceptual centrality in terms of dependency relations.

It is one thing to say that a hypothesis is not inconsistent with knowledge differences, and quite another to give a positive account of when and how knowledge affects concept use. Gentner and Medina (1998) in their discussion on adult induction draw an interesting parallel between the similarity versus rules distinction in induction, and Newell and Simon's (1972) distinction between *weak methods* and *strong methods*. Weak methods are general strategies (such as modus ponens or the premise diversity principle) that can operate under insufficient knowledge of a domain. Strong methods are methods that make heavy use of represented knowledge. Gentner and Medina believe that strong methods are more appropriate when sufficient knowledge is present. This dovetails nicely with the Lopez et al.'s account of their results: Americans relied on a weak method (the diversity principle) because they had insufficient knowledge, while Itzajs and tree experts relied heavily on strong methods (on represented knowledge, such as ecological knowledge) to make an inference.

8.6.5 Summary

Up to now we have presented evidence that demonstrated judgmental instabilities across tasks, development, and cultures and argued that (i) none of them individually is inconsistent with the notion of (fairly) stable mental representations, and (ii) that some differences (like cultural instabilities) sit well with our hypothesis since they are attributable to knowledge differences. However, some theorists considering some of this or other evidence simultaneously have reached the opposite conclusion that concepts are highly unstable entities (Barsalou, 1983, 1989; William James, 1950; Smith and Samuelson, 1997). Barsalou (1993) proposed that we should re-conceptualize 'concepts' as referring to momentary thoughts rather than timeless internal structures, while William James (1950) advocated that concepts are 'mythological entities'. It is outside the scope of the present discussion to fully expose the issue of stable versus unstable representations (for a review see Smith and Samuelson, 1997). We are content to state that the centrality hypothesis works even if the notion of stable representations is relaxed. Even if most of concepts are formed on-line, like Barsalou's (1983) *ad hoc* concepts, the notion of centrality would still be applicable

as long as such on-line entities are structured. There is every reason to believe that such entities will be constrained by knowledge or pre-theoretical biases and hence reasons to believe that the centrality hypothesis will be applicable.

8.7 CONCEPTUAL CENTRALITY AND ESSENTIALISM

The notion of conceptual centrality begs to be compared with the notion of essentialism. Ahn (1998) presents a very good comparison between the causal status hypothesis and various flavors of essentialism. The causal status hypothesis can be seen as a special case of conceptual centrality since (in its most relaxed form) it leaves out non-explanatory dependency relations like implications, or ones coming from pre-theoretical intuitions. What follows is a presentation of Ahn's arguments (with slight modifications) as applied to conceptual centrality.

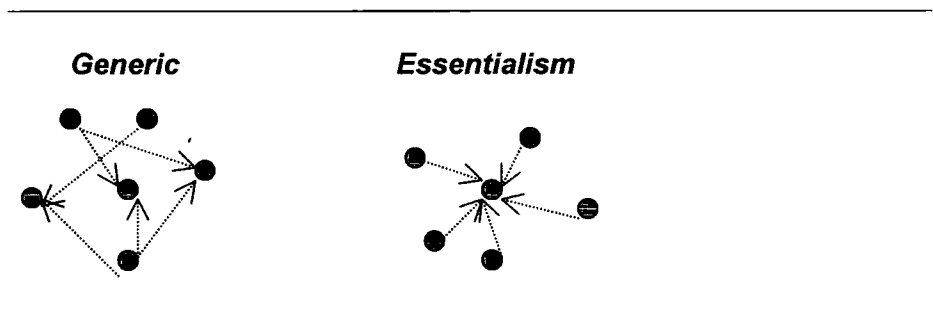
A pervasive philosophical and psychological view about concepts is that (at least some) things have essences. According to Locke (1894/1975) essences refer to unobservable properties of things that make them the thing they are. Putnam (1977) similarly advocates that natural kinds have essences although we may not know what their essences are. More recently Medin and Ortony (1989) proposed psychological essentialism: the position that irrespective of whether or not things really have essences, people act as if they do have essences that explain their phenomenal features.

The notion of conceptual centrality shares with essentialism the assumption that some features are more 'central' than others. Conceptual centrality, at the same time, differs from particular flavors of essentialism in that (1) it does not assume that central features are necessarily defining features, (2) it does not divide features into surface and deep features, (3) it is not restricted to natural kind categories like some essentialist accounts, (4) it is not independent upon the theories (knowledge) that a person has about a concept. Each of these four points is clarified below.

8.7.1 Central features are not necessarily defining features

According to a strong version of essentialism essences are defining features: if a thing has an essence of a certain kind then it is an object of that kind, whereas if the thing does not have an essence of a certain kind then it cannot be a member of that kind. According to conceptual centrality the features of a particular concept can be put on a continuum reflecting their relative immutability. Conceptual centrality is not committed as to whether essences (in our terms, features upon which all other features either directly or indirectly depend) in either a psychological or metaphysical sense exist. Figure 8.1 shows two possible dependency structures, a general one and one that reflects psychological essentialism (see Keil, 1995, p.237, for a variety of possible causal patterns).

Figure 8. 1 Simplified illustration of a generic dependency structure, and a special one reflecting essentialism. Dots stand for properties.



Keil (1995) presents arguments against metaphysical essentialism. He argues, for instance, that although some theorists equate the essence of animals with some DNA sequence, a DNA sequence cannot qualify as an essence since no two animals share the same sequence of nucleotides. Note that essentialism does not sit well with Darwinian evolution - in fact the great insight of natural selection emerged only when Wallace and Darwin abandoned the strongly held ideas of species having essences. Accepting that two contemporary species (presumably two different natural kinds) derive from a common ancestor entails accepting the splitting up of an essence.

As Medin and Ortony (1989) point out, as psychologists we should be more interested in what people believe about things rather than in about what things really are. Empirical evidence corroborates psychological essentialism. McNamara and Sternberg (1983) showed that many people believe that some properties are defining. Psychological essentialism seems to be universal. Atran (1987) showed psychological essentialism beliefs across cultures, and Gelman and Wellman (1991) found that such beliefs are present even in young children (see also Hatano, Siegler, Inagaki, Stavy, & Wax, 1993).

8.7.2 Central features are not necessarily deep features

Locke (1894/1975) believed that essences are undiscoverable. In a similar vein, some essentialists (e.g., Putnam, 1977) categorize features as either essential or surface ones. However, we believe that how deep a property actually is in a thing is orthogonal to its conceptual centrality. That is some deep properties are non-central (e.g. an unexpressed DNA sequence) and some surface properties are central (e.g. the color of a solar panel). To make a broad analogy, conceptual centrality is an increasing function of how 'deep' a feature is in the *dependency structure* (or intuitive theory) that an individual has about a concept - not in the actual thing itself. An unexpressed DNA sequence is a deep but nonetheless mutable property of an animal since almost nothing (according to most people's knowledge) depends on it. The color of a solar panel is one of its surface but nonetheless immutable properties since its crucial functions of attracting and storing energy depend on it. To summarize, we believe that features vary in conceptual centrality with features in the center of a dependency structure being more central than those in the periphery.

8.7.3 Conceptual centrality depends on (specific or vague) knowledge

According to some essentialist accounts (e.g. Kripke, 1971; Putnam, 1975) the categorization of natural kinds is independent of our knowledge of essential properties. Even if scientists discovered that gold is actually composed of plasticine, that would not change what we have been referring to as gold. On this account, what we refer as gold picks out some

stuff independently of how much (or if) we know about that stuff (see also Komatsu, 1992). On the other hand, the conceptual centrality of a feature directly derives from the theories (or vague ideas) that someone has about that concept. If our knowledge of a concept is changed, then its dependency structure and hence the centrality weights of some of its features might change as well. If scientists discover (and we get informed about this discovery) that sharks are not actually fish but are robots controlled by aliens and their appearance is but a disguise to fool us, then we might reconsider the centrality of their bronchia. In fact, not only specific knowledge influences conceptual centrality but even vague pre-theoretical biases.

8.7.4 Conceptual centrality is domain-independent

Although some theorists hold that both natural and nominal kinds have essences (e.g., Putnam, 1975), some others reserve essences to natural kinds (e.g., Schwartz, 1979). In general there are certainly many differences between artifact and animal categories (for a review see Keil, 1989; Keil et al. 1998). Nonetheless, we hold that conceptual centrality offers a domain-general mechanism of weighting features. What our account presupposes is knowledge (or at a minimum some intuition) of how properties of a kind are interconnected. To be concrete, assume that a person's representation of artifact categories are centered around their functional features whereas representations of natural kinds are centered around molecular features (see e.g. Barton & Komatsu, 1989; Gelman, 1988; Keil, 1989; Rips, 1989; but see Malt & Johnson, 1992). The conceptual centrality mechanism will not discriminate among these categories but it will promptly show functional features to be more central to artifacts, and molecular features to be more central to natural kinds. This is in fact what Ahn (1998) found. In essence, conceptual centrality is a structural constraint so differences in content or differences in clustering do not matter (see Figure 7.1).

8.8 THE PROBLEM OF INDUCTION

The current research addressed two issues about category-based inference: (1) Why are people more willing to generalize some properties over others? and (2) Why does the

generalizability of a property depend on the premise and conclusion categories? Our account of both issues is in terms of conceptually centrality. People are more willing to generalize some properties over others because certain properties promote more coherence to the target - they support more of its central properties. Furthermore, people's willingness to generalize a property depends on the specific concepts used because centrality is concept-specific - the same property may lend lots of support to certain targets but not to others. *Being round* is projectible for wheels but not for oranges, because roundness is central for understanding wheels but not oranges.

Have we solved the problem of induction then? No! We just have identified a strategy people use to project properties when (even vague) knowledge about their centrality status is available. Other strategies people use when relevant knowledge is available include causal, and gap-like strategies. Categorical arguments of the form *Lakes are contaminated with toxic waste; therefore fish in the lakes are contaminated with toxic waste* are strong because of a strong causal link between the argument's premise and conclusion. Arguments of the form *Poodles can bite through barbed wire; therefore German shepherds can bite through barbed wire* seem to invite gap-like reasoning. Future research should identify when each strategy is employed.

APPENDIX

**MATERIALS, ANALYSES OF VARIANCE
& CONTENTS OF THE
ACCOMPANYING FLOPPY DISC**

Table A.2.1 F_1 ANOVA of the centrality scores for Experiment 1.**Tests of Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	41.860	1	41.860	91.166	.000
Centr * Task order	.162	1	.162	.352	.560
Error(Centr)	9.183	20	.459		

Tests of Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Task order	14.972	1	14.972	2.064	.166
Error	145.060	20	7.253		

Table A.2.2 F_2 ANOVA of the centrality scores for Experiment 1.

Source	SS	df	MS	F	Sig.
Centrality	45.665	1	45.665	557.728	.000
Error(Centr)	.901	11	8.188E-02		
Task order	16.333	1	16.333	197.334	.000
Error(Task order)	.910	11	8.277E-02		
Centr * Task order	.176	1	.176	3.297	.097
Error(Cen * task)	.588	11	5.347E-02		

Table A.2.3 F_1 ANOVA of the scores from Experiment 1.

Source	SS	df	MS	F	Sig.
Order	8.437E-03	1	8.437E-03	.225	.640
Error	.824	22	3.744E-02		
Total	.832	23			

Table A.2.4 F_1 ANOVA of the centrality scores for Experiment 2**Tests of Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	120.447	1	120.447	263.416	.000
Cent * Cent Measure	1.098E-02	1	1.098E-02	.024	.878
Error(Centr)	19.662	43	.457		

Tests of Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Centr measure	30.525	1	30.525	9.449	.004
Error	138.905	43	3.230		

Table A.2.5 F_2 ANOVA of the centrality scores for Experiment 2

Source	SS	df	MS	F	Sig.
Centrality	.784	1	.784	7.519	.019
Error(Centr)	1.147	11	.104		
Centr measure	.102	1	.102	.526	.484
Error(centr measure)	2.126	11	.193		
Centr * Centr measure	.438	1	.438	5.615	.037
Error(Centr * Centr measure)	.858	11	7.798E-02		

Table A.2.6 F_1 ANOVA of the conditional likelihood scores from Experiment 2.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	1.995	1	1.995	6.848	.012
Centr * Centr_Def	.801	1	.801	2.749	.105
Error(Centr)	12.528	43	.291		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Centr_Def	.170	1	.170	.114	.737
Error	63.923	43	1.487		

Table A.2.7 F_2 ANOVA of the conditional likelihood scores from Experiment 2.

Source	SS	df	MS	F	Sig.
Centrality	.784	1	.784	7.519	.019
Error(Centrality)	1.147	11	.104		
Centr_Def	.102	1	.102	.526	.484
Error(Centr_def)	2.126	11	.193		
Centr * Centr_Def	.438	1	.438	5.615	.037
Error(Centr*Centr_Def)	.858	11	7.798E-02		

Table A.2.8 F_1 ANOVA of the unconditional likelihood scores from Experiment 2.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	7.813	1	7.813	13.577	.001
Centr * Centr_Def	1.444	1	1.444	2.510	.120
Error(Centr)	24.743	43	.575		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Centr_Def	.614	1	.614	1.269	.266
Error	20.820	43	.484		

Table A.2.9 F_2 ANOVA of the unconditional likelihood scores from Experiment 2.

Source	SS	df	MS	F	Sig.
Centrality	3.731	1	3.731	17.431	.002
Error(Centrality)	2.354	11	.214		
Centr_Def	.350	1	.350	3.709	.080
Error(Centr_Def)	1.039	11	9.446E-02		
Centr * Centr_Def	1.175	1	1.175	3.942	.073
Error(Centr*Centr_Def)	3.277	11	.298		

Table A.3.1 The argument premises (facts) of Experiments 3 and 4. Control items in both experiments contained only the description in bold-faced letters. Causal+explanation items (Experiment 3) contained the whole description. Causal items (Experiment 4) contained the whole description except the text inside brackets. Only one position of properties (the one where the central property comes first) is shown.

Argument premises (Facts)	Conclusion
<p>Cows have acute kinaesthetic receptors and good hand eye co-ordination. For cows, acute kinaesthetic receptors cause good hand-eye co-ordination [because higher controllability of motor movements promotes good co-ordination.]</p>	<u>High Sim</u> Horse
<p>Squirrels have exceptional hearing and frequent headaches. For squirrels, exceptional hearing causes frequent headaches [because very high auditory frequencies can cause nasal cavities to vibrate.]</p>	Mouse
<p>Seals have a rapid heart rate and a high water consumption. For seals, rapid heart rate causes high water consumption [because it lowers blood pressure which can be corrected by high consumption of liquids.]</p>	Dolphin
<p>Elephants have a strong odour and restless sleeping habits. For elephants, a strong odour causes restless sleeping habits [because it induces high amphetamine production that results in sleeplessness.]</p>	Rhino
<p>Mice have high amounts of oxydilic acid and good digestion. For mice, high amounts of oxydilic acid cause good digestion [because they help in the breaking down of fat.]</p>	Squirrel
<p>Dolphins have vitamin deficiency and overstretched vocal chords. For dolphins, vitamin deficiency causes overstretched vocal chords [because deficiency of essential vitamins affect the normal development of the larynx.]</p>	Seal
<p>Rhinos have high amounts of opiates and accident susceptibility. For rhinos, high amounts of opiates cause accident susceptibility [because they induce undersensitive pain receptors that inhibit species-typical defensive responses.]</p>	Elephant
<p>Horses have a weak immune system causes and an acute sense of smell. For horses, a weak immune system causes an acute sense of smell [because it makes the animal susceptible to food poisoning which can be avoided with sensitive smell detectors.]</p>	Cow

Table A.3.1 (continued)

	<u>Low-sim</u>
<p>Cows have high amounts of glucose and an increased metabolism. For cows, high amounts of glucose cause an increased metabolism [because they stimulate glutamate accumulation into the cell and its oxidation.]</p>	Dolphin
<p>Squirrels have an underdeveloped hippocampus and anterograde amnesia. For squirrels, an underdeveloped hippocampus causes anterograde amnesia [because underdeveloped mammillary bodies aversively affect memory.]</p>	Rhino
<p>Seals have low REM sleep and severe learning impairments. For seals, low REM sleep causes severe learning impairments [because low REM sleep impairs the ability to concentrate.]</p>	Horse
<p>Elephants have dense hair and little skin pigment. For elephants, dense hair permits little skin pigment [because the hair can cover the body and protect it from the sun, a function usually served by pigment.]</p>	Mouse
<p>Mice have high amounts of binaural neurons and good localisation of stimuli. For mice, high amounts of binaural neurons cause good localisation of stimuli [because they respond to changes in arrival times of sounds presented to both ears.]</p>	Elephant
<p>Dolphins have an overactive thyroid gland and high amounts of peptides. For dolphins, an overactive thyroid gland permits high amounts of peptides [because it leads to increased metabolism which compensates for high food consumption.]</p>	Cow
<p>Rhinos have low amounts of ADH hormone and high osmotic pressure. For rhinos, low amounts of ADH hormone cause high osmotic pressure [because they regulate the amount of water that the body excretes.]</p>	Squirrel
<p>Horses have a loose tensor tympani and frequent migraines. For horses, a loose tensor tympani results in frequent migraines [because it regulates the amount of sound that is permitted to pass through the ear.]</p>	Seal

Table A.3.2 F_1 ANOVA of the scores from Experiment 3.

Source	SS	df	MS	F	Sig.
Shared structure	7.878E-02	1	7.878E-02	1.208	.283
Error(Shared)	1.499	23	6.519E-02		
Relation	.287	1	.287	4.285	.050
Error(Relation)	1.541	23	6.700E-02		
Shared * Relation	1.628E-02	1	1.628E-02	.240	.629
Error(Shared*Relation)	1.562	23	6.791E-02		

Table A.3.3 F_2 ANOVA of the scores from Experiment 3.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Relation	9.790E-02	1	9.790E-02	13.241	.003
Relation * Shared	6.328E-03	1	6.328E-03	.856	.371
Error(relation)	.104	14	7.394E-03		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Shared structure	2.820E-02	1	2.820E-02	.485	.497
Error	.814	14	5.811E-02		

Table A.3.4 F_1 ANOVA of the scores from Experiment 4.

Source	SS	df	MS	F	Sig.
Relation	.510	1	.510	6.192	.021
Error(Relation)	1.896	23	8.243E-02		
Shared Structure	.315	1	.315	8.616	.007
Error(Shared)	.841	23	3.657E-02		
Relation * Shared	.000	1	.000	.000	1.000
Error(Relat*Shared)	.969	23	4.212E-02		

Table A.3.5 F_2 ANOVA of the scores from Experiment 4.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Relation	.170	1	.170	15.955	.001
Relat * Shared	3.125E-10	1	3.125E-10	.000	1.000
Error(Relation)	.149	14	1.067E-02		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Shared Structure	.105	1	.105	1.739	.208
Error	.846	14	6.041E-02		

Table A.3.6 Descriptions of the 18 pairs of properties used in Experiments 6 and 7. Central properties are presented first. Semantic relations are highlighted.

Triple 1

The enzyme lipase regulates metabolism.
The enzyme protease is regulated by metabolic rate.

The hormone prolactin regulates blood flow.
The hormone renin is regulated by blood flow.

The neurotransmitter acetylcholine helps detect predators.
The levels of the neurotransmitter nonadrenaline increase after seeing a predator.

Triple 2

The enzyme amylase helps dissolve blood clots.
The enzyme streptokinase gets denser when blood clots.

The hormone ACTH modulates blood pressure.
The hormone LH is regulated by blood pressure.

The neurotransmitter GABA helps oxygenate the brain.
The neurotransmitter glycine depends on the oxygen level of the brain.

Triple 3

The hormone TSH ensures normal muscle development.
The hormone FSH is produced as a result of muscle development.

The neurotransmitter octopamine determines food intake.
The production of the neurotransmitter histamine is determined by food intake.

The enzyme papain reduces inflammation.
The enzyme trypsin can double in size because of inflammation.

Triple 4

The hormone ADH helps the animal cope with stress.
The hormone MSH is produced in response to stress.

The neurotransmitter norepinephrine regulates movement.
The neurotransmitter epinephrine is regulated by their movements.

The enzyme thrombin helps detect poisonous food.
The enzyme subtilisin changes color in response to poisonous food.

Triple 5

The neurotransmitter tyrosine regulates many brain functions.
The neurotransmitter taurine is regulated by many brain functions.

The enzyme elastase helps to digest food.
The enzyme aliesterase changes shape during digestion.

The hormone oxytocin regulates mating behavior.
The levels of the hormone calcitocin increase after mating.

Triple 6

The neurotransmitter quisqualate promotes visual acuity.
The neurotransmitter kainate depends on visual acuity.

The enzyme phosphorylase regulates blood pH.
The levels of the enzyme transacetylase are regulated by blood pH.

The hormone vassopresin controls bone growth.
The hormone proctolin is secreted as a result of bone growth.

Table A.3.7 F_1 ANOVA of the scores from Experiment 6.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.666	30.909	2.000	31.000	.000
Centrality	Pillai's Trace	.129	4.739	1.000	32.000	.037
Shared * Centr.	Pillai's Trace	.046	.750	2.000	31.000	.481

Table A.3.8 F_2 ANOVA of the scores from Experiment 6.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	26.694	1	26.694	12.416	.003
Centr * Shared	1.556	2	.778	.362	.702
Error(Centr)	32.250	15	2.150		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Shared Structure	3844.667	2	1922.333	27.737	.000
Error	1039.583	15	69.306		

Table A.3.9 F_1 ANOVA of the scores from Experiment 7.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.750	46.501	2.000	31.000	.000
Centrality	Pillai's Trace	.129	4.748	1.000	32.000	.037
Shared * Centrality	Pillai's Trace	.046	.752	2.000	31.000	.480

Table A.3.10 F_2 ANOVA of the scores from Experiment 7.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	191.890	1	191.890	19.553	.000
Central * Shared	7.233	2	3.616	.368	.698
Error(Central)	147.209	15	9.814		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Shared Structure	5841.360	2	2920.680	42.157	.000
Error	1039.206	15	69.280		

Table A.4.1 F_1 ANOVA of the scores from Experiment 8.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Shared Structure	3.126	1	3.126	30.445	.000
Shared * List	2.083E-02	1	2.083E-02	.203	.657
Error(Shared)	2.259	22	.103		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
List	5.208E-03	1	5.208E-03	.048	.828
Error	2.368	22	.108		

Table A.4.2 F_2 ANOVA of the scores from Experiment 8.

Source	SS	df	MS	F	Sig.
Shared Structure	1.021	1	1.021	275.577	.000
Error	5.187E-02	14	3.705E-03		

Table A.4.3 F_1 ANOVA of the scores from Experiment 9.

Source	SS	df	MS	F	Sig.
Shared Structure	.647	1	.647	9.758	.005
Error(shared)	1.525	23	6.628E-02		
Centrality	.258	1	.258	7.823	.010
Error(Centr)	.760	23	3.302E-02		
Shared * Centr	.252	1	.252	5.281	.031
Error(Shared*Centr)	1.098	23	4.775E-02		

Table A.4.4 F₁ ANOVA of the scores from Experiment 10.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared Structure	Pillai's Trace	.456	9.205	2.000	22.000	.001
Centrality	Pillai's Trace	.334	11.511	1.000	23.000	.003
Shared * Centrality	Pillai's Trace	.339	5.646	2.000	22.000	.010

Table A.4.5 F₂ ANOVA of the scores from Experiment 10.

Within-Subjects Effects

Source	SS	df	MS	F	Sig.
Centrality	4.988E-02	1	4.988E-02	24.226	.000
Centr * Shared	8.544E-02	2	4.272E-02	20.749	.000
Error(Centr)	3.088E-02	15	2.059E-03		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Shared Structure	.245	2	.122	50.844	.000
Error	3.608E-02	15	2.406E-03		

Table A.4.6 F₁ ANOVA of the scores from Experiment 11.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared Structure	Pillai's Trace	.446	12.877	2.000	32.000	.000
Centrality	Pillai's Trace	.366	9.247	2.000	32.000	.001
Shared * Centr	Pillai's Trace	.424	5.530	4.000	30.000	.002

Table A.4.7 F_1 ANOVA (Shared structure) x (Centrality: central or less-central) of the scores from Experiment 11.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.498	15.869	2.000	32.000	.000
Centrality	Pillai's Trace	.312	14.982	1.000	33.000	.000
Shared * Centrality	Pillai's Trace	.277	6.145	2.000	32.000	.006

Table A.4.8 F_1 ANOVA (Shared structure) x (Centrality: central or no-information given) of the scores from Experiment 11.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.482	14.907	2.000	32.000	.000
Centrality	Pillai's Trace	.365	18.976	1.000	33.000	.000
Shared * Centrality	Pillai's Trace	.412	11.221	2.000	32.000	.000

Table A.5.1 F_1 ANOVA of the frequency scores for the novel properties from Experiment 12.

Within-Subjects Effects

Source	SS	df	MS	F	Sig.
Centrality	17.521	1	17.521	8.769	.007
Central * Order	11.021	1	11.021	5.516	.028
Error(Central)	43.958	22	1.998		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Order	.521	1	.521	.029	.867
Error	397.958	22	18.089		

Table A.5.2 F_1 ANOVA of the conditional likelihood scores for the novel properties from Experiment 12.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.484	6.577	3.000	21.000	.003
Centrality	Pillai's Trace	.358	12.805	1.000	23.000	.002
Shared * Centr	Pillai's Trace	.059	.437	3.000	21.000	.729

Table A.5.3 F_1 ANOVA of the frequency scores for the novel properties from Experiment 13.

Within-Subjects Effects

Source	SS	df	MS	F	Sig.
Centrality	1.687	1	1.687	.250	.622
Central * Order	15.187	1	15.187	2.248	.148
Error(Central)	148.625	22	6.756		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Order	35.021	1	35.021	1.830	.190
Error	420.958	22	19.134		

Table A.5.4 F_1 ANOVA of the conditional likelihood scores for the novel properties from Experiment 13.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.409	4.835	3.000	21.000	.010
Centrality	Pillai's Trace	.387	14.541	1.000	23.000	.001
Shared * Centr	Pillai's Trace	.434	5.369	3.000	21.000	.007

Table A.5.5 F_1 ANOVA (Centrality) x (Shared structure: rhinos or mice or falcons) of the conditional likelihood scores for the novel properties from Experiment 13.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.352	5.967	2.000	22.000	.009
Centrality	Pillai's Trace	.442	18.194	1.000	23.000	.000
Shared * Centrality	Pillai's Trace	.200	2.751	2.000	22.000	.086

Table A.5.6 F_1 ANOVA of the frequency scores for the novel properties from Experiment 14

Within-Subjects Effects

Source	SS	df	MS	F	Sig.
Centrality	2.970E-02	1	2.970E-02	.033	.856
Centr * Order	.644	1	.644	.719	.401
Centr * Freq. cond.	.377	2	.188	.210	.811
Centr * order * Freq. cond.	.377	2	.188	.210	.811
Error(Centr)	41.230	46	.896		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Order.	99.704	1	99.704	3.738	.059
Freq. cond.	46.293	2	23.147	.868	.427
Order * Freq. Cond.	94.869	2	47.434	1.779	.180
Error	1226.855	46	26.671		

Table A.5.7 F_1 ANOVA of the conditional likelihood scores for the novel properties from Experiment 14.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.309	7.297	3.000	49.000	.000
Centrality	Pillai's Trace	.279	19.765	1.000	51.000	.000
Shared * Centrality	Pillai's Trace	.241	5.200	3.000	49.000	.003

Table A.5.8 F_1 ANOVA (Centrality) x (Shared structure: rhinos or mice or falcons) of the conditional likelihood scores for the novel properties from Experiment 14.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.280	9.698	2.000	50.000	.000
Centrality	Pillai's Trace	.274	19.258	1.000	51.000	.000
Shared * Centrality	Pillai's Trace	.078	2.101	2.000	50.000	.133

Table A.5.9 F_1 ANOVA of the frequency scores for the novel properties from Experiment 14

Within-Subjects Effects

Source	SS	df	MS	F	Sig.
Centrality	23.784	1	23.784	4.778	.034
Centr * Freq. cond.	3.045	2	1.522	.306	.738
Centr * Order	2.131	1	2.131	.428	.516
Centr * Freq. cond. * Order	6.319	2	3.160	.635	.535
Error(Centr)	228.989	46	4.978		

Table A.5.10 F_1 ANOVA of the conditional likelihood scores for the familiar properties from Experiment 14.

Effect		Value	F	Hypothesis df	Error df	Sig.
Centrality	Pillai's Trace	.690	113.663	1.000	51.000	.000
Shared Structure	Pillai's Trace	.840	86.049	3.000	49.000	.000
Centr * Shared	Pillai's Trace	.718	41.635	3.000	49.000	.000

Table A.6.1 F_1 ANOVA of the scores from Experiment 15

Within-Subjects Effects

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.489	9.561	2.000	20.000	.001
Shared * List	Pillai's Trace	.160	.914	4.000	42.000	.465

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
List	23.250	2	11.625	1.732	.201
Error	140.917	21	6.710		

Table A.6.2 F_2 ANOVA of the scores from Experiment 15.

Source	SS	df	MS	F	Sig.
Shared structure	291.000	2	145.500	80.833	.000
Error	27.000	15	1.800		
Total	4736.000	18			

Table A.6.3 Properties used in Experiment 16.

Triple #	Property Names		
	HSSS	LSSS	LSDS
1	Tismine	Pirtine	warkine
2	Terone	Maldone	quazone
3	Perrain	Trypain	strylain
4	Trolone	Panctone	zylone
5	Fermain	Saltain	papain
6	Trypsine	Kermine	sufine

Table A.6.4 F_1 ANOVA of the scores from Experiment 16.

Effect		Value	F	Hypothesis df	Error df	Sig.
Shared structure	Pillai's Trace	.113	1.845	2.000	29.000	.176

Table A.6.5 F_2 ANOVA of the scores from Experiment 16.

Source	SS	df	MS	F	Sig.
Shared structure	68.778	2	34.389	23.808	.000
Error	21.667	15	1.444		
Total	10078.000	18			

Table A.7.1 F_1 ANOVA of the scores from Experiment 17.

Effect		Value	F	Hypothesis df	Error df	Sig.
Centrality	Pillai's Trace	.381	12.911	1.000	21.000	.002
Shared	Pillai's Trace	.855	58.795	2.000	20.000	.000
Centr * Shared	Pillai's Trace	.459	8.476	2.000	20.000	.002

Table A.7.2 F_2 ANOVA of the scores from Experiment 17.**Within-Subjects Effects**

Source	SS	df	MS	F	Sig.
Centrality	943.862	1	943.862	34.826	.000
Centr * Shared	619.316	2	309.658	11.426	.000
Error(Centr)	569.140	21	27.102		

Between-Subjects Effects

Source	SS	df	MS	F	Sig.
Shared structure	4185.873	2	2092.937	17.559	.000
Error	2503.118	21	119.196		

Table A.8.1 Mean (SE) Centrality and inductive strength estimates for study on wingless Falcons versus wingless chicken.

Centrality Task	
Chickens	5.13 (.32)
Falcons	5.70 (.35)
Induction Task	
Chickens	6.17 (1.94)
Falcons	2.77 (1.45)

Note. In both tasks a 0 to 100 scale was used. For the centrality task, higher estimates indicate higher centrality ratings for the property *wings*. For the induction task, higher estimates indicate higher likelihood ratings for the possibility of a *wingless* species. Wings were rate marginally more central for falcons than for chicken ($t(36)=1.95, p<.06$). The possibility of wingless falcons was rated significantly less likely than the possibility of wingless chickens ($t(36)= -2.32, p<.05$).

Table A.8.2 Proof that a negative centrality model can capture negative centrality effects.

Assume that,

The overlap between the base and the target with respect to the dependent properties (BandT, d-on) is constant a , where $0 < a < 1$.

So, amount of non-overlap (Band~T, d-on) will be $1-a$, where $0 < 1-a < 1$

Say c^+ and c^- are two non-negative constants that stand for the number of dependent properties of the candidate feature in the base. In a sense, these reflect the maximum number of properties the target can share with the base. $c^+ - c^- > 0$ because, by definition, more properties depend on the central feature.

Then, the effect of centrality for the central property is, $a c^+ - (1-a) c^+$
and the effect of centrality for the less central property is, $a c^- - (1-a) c^-$

So, the centrality effect is,
 $[a c^+ - (1-a) c^+] - [a c^- - (1-a) c^-] = a (c^+ - c^-) - (1-a) (c^+ - c^-)$

CE > 0, when $a (c^+ - c^-) > (1-a) (c^+ - c^-)$
when $a > 1-a$, (dividing by common factor) [$c^+ - c^- > 0$, assumption]
when $.5 < a < 1$.

So, CE < 0 when $0 < a < .5$

Contents of the accompanying floppy disk

The floppy disk accompanying this thesis contains the following 15 MS Excel files (compatible with version MS Excel 97 or MS Excel 5.0/95 for Windows).

File Name

Chapter II, Experiment 1
Chapter II, Experiment 2
Chapter III, Experiments 3 & 4
Chapter III, Experiments 6 & 7
Chapter IV, Experiment 8
Chapter VI, Experiment 9
Chapter VI, Experiment 10
Chapter VI, Experiment 11
Chapter V, Experiment 12
Chapter V, Experiment 13
Chapter V, Experiment 14
Chapter VI, Experiment 15
Chapter VI, Experiment 16
Chapter VII, Experiment 17
Similarity data for Experiments 6, 7, 10, & 11.

Each of these files contains information that produced the statistics connected with one or more Experiments. The information of each file is self explanatory. The file's name indicates the Chapter, and the Experiment number (e.g. "Chapter II, Experiment 1"). The sheet labels within a file (presented at the bottom of each Excel sheet) indicate how the data were organized (e.g. "DATA FOR Ss"). The variables within each file are illustrated clearly, and legends are available to further aid clarification.

BIBLIOGRAPHY

- Adamson, R. E., 1952. Functional fixedness as related to problem solving: A repetition of three experiments, *Journal of Experimental Psychology*, 44, 288-291.
- Ahn, W. (1998). Why are different features central for natural kinds and artifacts?: the role of causal status in determining feature centrality. *Cognition*, 69, 135-178.
- Ahn, W., & Dennis, M. J. (1997). Dissociation between categorization and similarity judgment: Differential effect of causal status on feature weights. *Proceedings of the Interdisciplinary Workshop on Similarity and Categorisation*, Edinburgh, Scotland.
- Ahn, W., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, 54, 299-352.
- Ahn, W., and Lassaline, M.E. (1995). Causal structure in categorization. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, Pittsburgh, PA, 521-526.
- Ahn, W., and Sloman, S. A. (1997). Distinguishing name centrality from conceptual centrality. *Proceedings of the Nineteenth Conference of the Cognitive Science Society*, Stanford, CA, 1-6.
- Atran, S. (1990). *Cognitive foundations of natural history: towards an anthropology of science*. Cambridge University Press.
- Atran, S. (1995). Causal constraints on categories and categorical constraints on biological reasoning across cultures. In D. Sperber, D. Premack, and A. J. Premack (Eds.), *Causal cognition: A multidisciplinary approach* (pp. 205-233). New York: Oxford University Press.
- Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavior and Brain Sciences*, 21, 547-569.
- Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infant's ability to draw inferences about non-obvious object properties: Evidence from exploratory play. *Child Development*, 64, 711-728.

- Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory and Cognition*, *15*, 397-418.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*, 211-227.
- Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- Barsalou, L. W. (1989). Intraconcept similarity and its implications for interconcept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76-121). New York: Cambridge University Press.
- Barton, M. E., & Komatsu, L. K. (1989). Defining features of natural kinds and artifacts. *Journal of Psycholinguistic Research*, *18*, 433-447.
- Birch, H. G., and Rabinowitz, H. S, 1951. The negative effect of previous experience in productive thinking, *Journal of Experimental Psychology*, *41*, 121-125.
- Carey, S (1985). *Conceptual change in childhood*. Cambridge, MA: Bradford Books.
- Carey, S. (1995). On the origins of causal understanding. In D. Sperber, D. Premack, and A. J. Premack (Eds.), *Causal cognition: A multidisciplinary approach* (pp. 268-302). New York: Oxford University Press.
- Chapman, L. J., & Chapman, J. P. (1967). Genesis of popular but erroneous diagnostic observations. *Journal of Abnormal Psychology*, *72*, 193-204.
- Chater, N., & Oaksford, M. (1993). Logicism, mental models, and everyday reasoning: reply to Garnham. *Mind and Language*, *8*, 72-89.
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, *58*, 545-567.
- Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review*, *99*, 365-382.

- Clark, H. H. (1973). The language as fixed effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behaviour*, 12, 335-359.
- Clement, J., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, 15, 89-132.
- Coley, J. D., Medin, D. L., & Atran, S. (1997). Does rank have its privilege? Inductive inferences within folkbiological taxonomies. *Cognition*, 64, 73-112.
- Cooper, J., & Fazio, R. H. (1984). A new look at dissonance theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 229-267). New York: Academic Press.
- Davidson, N. S. & Gelman, S. A. (1990). Inductions from novel categories: The role of language and conceptual structure. *Cognitive Development*, 5, 151-172.
- Dawkins, R. (1986). *The blind watchmaker*. Great Britain: Bath Press.
- Duncker, K., 1945. On problem-solving, *Psychological Monogram*, 58, no. 5 (whole no. 270).
- Dupre, J. (1981). Biological taxa as natural kinds. *Philosophical Review*, 90, 66-90.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Franks, B. (1995). Sense generation: A "quasi-classical" approach to concepts and concept combination. *Cognitive Science*, 19, 441-505.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Gelman, S. A. (1988). The development of induction with natural kind and artifact categories. *Cognitive Psychology*, 20, 65-95.

- Gelman, S. A., Collman, P., & Maccoby, E.E. (1986). Inferring properties from categories versus inferring categories from properties: The case of gender. *Child Development*, 57, 396-404.
- Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, 23, 183-209.
- Gelman, S. A., & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*, 58, 1532-1541.
- Gelman, S. A., & Wellman, H.M. (1991). Insides and essences: early understandings of the nonobvious. *Cognition*, 38, 213-144.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). New York: Cambridge University Press.
- Gentner, D., & Clement J. (1998). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp.307-358). New York: Academic Press.
- Gentner, D., & Medina, J. (1998). Similarity and the development of rules In S. A. Sloman & L. J. Rips (Eds.), *Similarity and symbols in human thinking* (pp. 177-211). Cambridge, MA: MIT Press.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 277-300.
- Goldstone, R. L. (1994a). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 3-28.
- Goldstone, R. L. (1994b). The role of similarity in categorization: providing a groundwork. *Cognition*, 52, 125-157.

- Goldstone, R. L., & Medin, D. L. (1994a). Similarity, interactive activation, and mapping: An overview. In K. Holyoak & J. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol.2: Analogical connections* (pp. 321-362). Norwood, New Jersey: Ablex.
- Goldstone, R. L., & Medin, D. L. (1994b). Time course comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 29-50.
- Goldstone, R. L., Medin, D. L., & Genter, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, *23*, 222-262.
- Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.
- Goodman, N. (1972). *Problems and projects*. New York: Bobbs-Merrill.
- Gutheil, G., & Gelman, S. A. (1997). Children's use of sample size and diversity information within basic-level categories. *Journal of Experimental Child Psychology*, *64*, 159-174.
- Hahn, U. & Chater, N. (1997). Concepts and similarity. In K. Lamberts and D. Shanks (Eds.), *Knowledge, Concepts, and Categories* (pp. 43-92). London: Psychology Press.
- Hampton, J. A. (1982). A demonstration of intransitivity in natural categories. *Cognition*, *12*, 151-164.
- Hampton, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory and Cognition*, *15*, 55-71.
- Hampton, J. A. (1997). Conceptual combination. In K. Lamberts and D. Shanks (Eds.), *Knowledge, Concepts, and Categories* (pp. 133-159). London: Psychology Press.
- Harman, G. (1995). Rationality. In D. Osherson and E.E. Smith (Eds.), *An invitation to cognitive science vol. 3: Thinking*. London: MIT Press.
- Hatano, G., Siegler, R. S., Inagaki, K., Stavy, R., & Wax, N. (1993). The development of biological knowledge: a multinational study. *Cognitive Development*, *8*, 47-62.

- Heit, E. (1997). Knowledge and concept learning. In K. Lamberts and D. Shanks (Eds.), *Knowledge, Concepts, and Categories* (pp. 7-41). London: Psychology Press.
- Heit, E. (in press). Properties of inductive reasoning. To appear in *Psychonomic Bulletin & Review*.
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 411-422.
- Hirschfield, L. A., & Gelman, S. A. (1994) (Eds.). *Mapping the mind*. Cambridge: Cambridge University Press.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: Bradford Books/MIT Press.
- Holyoak, K. J., & Thagard, P. R. (1989a). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Holyoak, K. J., & Thagard, P. R. (1989b). A computational model of analogical problem solving. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 242-266). New York: Cambridge University Press.
- Howell, D. C. (1997). *Statistical methods for psychology*. USA: Duxbury Press.
- Hume, D. (1748-1988). *An inquiry concerning human understanding*. Illinois: Open Court.
- James, W. (1950). *The principles of psychology*, vol 1. New York: Dover Publications. [Original work published in 1890. New York: Henry Holt & Co.]
- Kahneman, D., & Miller, D. T. (1986). Norm theory: comparing reality to its alternatives. *Psychological Review*, 93, 136-153.
- Kalish, C. W., & Gelman, S. A. (1992). On wooden pillows: Multiple classifications and children's category-based inductions. *Child Development*, 63, 1536-1557.
- Keil, F. C. (1989). *Concepts, Kinds, and cognitive development*. Cambridge, MA: MIT Press.

- Keil, F. C. (1995). The growth of causal understanding of natural kinds. In D. Sperber, D. Premack, and A. J. Premack (Eds.), *Causal cognition: A multidisciplinary approach* (pp. 234-262). New York: Oxford University Press.
- Keil, F. C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior*, 23, 221-236.
- Keil, F. C., Smith, C. W., Simons D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. In S. A. Sloman & L. J. Rips (Eds.), *Similarity and symbols in human thinking* (pp. 17-49). Cambridge, MA: MIT Press.
- Kelley, H. H. (1967). Attribution theory in social psychology. *Nebraska Symposium on Motivation* 15, 192-238.
- Kelley, H. H. (1973). The process of causal attribution. *American Psychologist*, 28, 107-128.
- Kempton, W. (1978). Category grading and taxonomic relations: a mug is a sort of cup. *American Ethnologist*, 5, 44-65.
- Komatsu, L. K. (1992). Recent views on conceptual structure. *Psychological Bulletin*, 112, 500-526.
- Kripke, S. (1972). Naming and necessity. In D. Davidson and G. Harman (Eds.), *Semantics of natural language* (2nd ed., pp.253-355). Dordrecht, Holland: Reidel.
- Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 754-770.
- Locke, J. (1894/1975). *An essay concerning human understanding*. Oxford: Oxford University Press.
- Lopez, A., Atran, S., Coley, J. D., Medin, D. L., & Smith, E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology*, 32, 251-295.

- Lopez, A., Gelman, S. A., Gutheil, G., & Smith, E. E. (1992). The development of category-based induction. *Child Development, 63*, 1070-1090.
- Lorch, R. F. Jr. & Myers, J. L. (1990). Regression analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 149-157.
- Maier, N. R. F., 1931. Reasoning in humans: The solution of a problem and its appearance in consciousness, *Journal of Comparative and Physiological Psychology, 11*, 181-194.
- Malt, B. C. (1994). Water is not H₂O. *Cognitive Psychology, 27*, 41-70.
- Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23*, 250-269.
- Malt, B. C., & Johnson, E. C. (1992). Do artifact concepts have cores? *Journal of Memory and Language, 31*, 195-217.
- Mandler, J. M., & McDonough, L. (1996). Drinking and driving don't mix: Inductive generalization in infancy. *Cognition, 59*, 307-335.
- Mandler, J. M., & McDonough, L. (1998). Studies in inductive inference in infancy. *Cognitive Psychology, 37*, 60-96.
- Markman, E. M. (1989). *Categorization and naming in children*. Cambridge, MA: MIT Press.
- Markman, A., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language, 32*, 517-535.
- Markman, A., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology, 25*, 431-467.
- McDonald, J., Samuels, M., & Rispoli, J. (1996). A hypothesis-assessment model of categorical argument strength. *Cognition, 59*, 199-217.
- McNamara, T. P., & Sternberg, R. (1983). Mental models of word meaning. *Journal of Verbal Learning and Verbal Behavior, 22*, 449-474.

- Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology*, 32, 49-96.
- Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179-196). New York: Cambridge University Press.
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, 20, 158-190.
- Medin, D. L., & Wattenmaker, W. D. (1987). Category cohesiveness, theories, and cognitive archeology. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 25-62). New York: Cambridge University Press.
- Mill, J. S. (1972) *A system of logic ratiocinative and inductive*. Toronto, Ontario, Canada: University of Toronto Press. (Originally published, 1843).
- Murphy, G. L. (1993). Theories and concept formation. In I. Van Mechelen, J. Hampton, R. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 173-200). London: Academic 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, 904-919.
- Murphy, G. L., & Kaplan, A. S. (1997). Category learning with minimal knowledge. *Proceedings of the Interdisciplinary Workshop on Similarity and Categorisation*, Edinburgh, Scotland.
- Murphy, G. L., and Lassaline, M. E. (1997). Hierarchical structure in concepts and the basic level of categorization. In K. Lamberts and D. Shanks (Eds.), *Knowledge, Concepts, and Categories* (pp. 93-131). London: Psychology Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.

- Murphy, G. L., & Wisniewski, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol. 2. Theory and applications* (pp. 23-45). Chichester: Ellis Horwood.
- Newell, A., & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.
- Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, *90*, 339-363.
- Oaksford, M., & Chater, N. (1991). Against logicist cognitive science. *Mind and Language*, *6*, 1-38.
- Osherson, D. N., Smith, E. E., Meyers, T. S., Shafir, E., & Stob, M. (1994). Extrapolating human probability judgment. *Theory and Decision*, *36*, 103-129.
- Osherson, D. N., Smith, E. E., & Shafir, E. B. (1986). Some origins of belief. *Cognition*, *24*, 197-224.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185-200.
- Osherson, D. N., Stern, J., Wilkie, O., Stob, M., & Smith, E. E. (1991). Default probability. *Cognitive Science*, *15*, 251-269.
- Pazzani, M. J. (1991). Influence of prior knowledge on concept acquisition: experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 416-432.
- Pennington, N. & Hastie, R. (1992). Reasoning in explanation based decision-making. *Cognition*, *49*, 123-163.
- Putnam, H. (1975). The meaning of "meaning." In K. Gunderson (Ed.). *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, VII, Minneapolis: University of Minnesota Press.

- Putnam, H. (1977). Is semantics possible? In Schwartz, S. P. (Ed.), *Naming, necessity, and natural kinds*. Ithaca, NY: Cornell University Press.
- Quine, W. V. O. (1961). Two dogmas of empiricism. In W. V. O. Quine, *From a logical point of view* (2nd ed., pp. Pp.20-46). New York: Harper and Row.
- Quine, W. V. O. (1977). Natural Kinds. In S. P. Schwartz (Ed.), *Naming, necessity, and natural kinds* (pp. 155-175). Ithaca, New York: Cornell University Press.
- Randall, R. A. (1976). How tall is a taxonomic tree? Some evidence for dwarfism. *American Ethnologist*, 3, 543-553.
- Rips, L. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning, and Verbal Behavior*, 14, 665-681.
- Rips, L. (1989). Similarity, typicality and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp.21-59). New York: Cambridge University Press.
- Rips, L., & Collins, A. (1993). Categories and resemblance. *Journal of Experimental Psychology: General*, 122, 468-486.
- Rosch, E., & Mervis, C. B. (1975) Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
- Ross, B. H., & Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology*, 38, 495-553.
- Shafir, E., Smith, E. E., & Osherson, D. (1990). Typicality and reasoning fallacies. *Memory and Cognition*, 18, 229-239.
- Simons, D. J., & Keil, F. C. (1995). An abstract to concrete shift in the development of biological thought: the insides story. *Cognition*, 56, 129-163.
- Skyrms, B. (1975). *Choice and chance: An introduction to inductive logic*. Encino, CA: Dickenson. 2nd ed.

- Sloman, S. A. (1993). Feature based induction. *Cognitive Psychology*, 25, 231-280.
- Sloman, S. A. (1994). When explanations compete: The role of explanatory coherence on judgments of likelihood. *Cognition*, 52, 1-21.
- Sloman, S. A. (1997). Explanatory coherence and the induction of properties. *Thinking and Reasoning*, 3, 81-110.
- Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology*, 35, 1-33.
- Sloman, S.A., Love B.C., and Ahn, W. (1998). Feature, centrality and conceptual coherence. *Cognitive Science*, 22, 189-228.
- Sloman, S. A., & Wisniewski, E. (1992). Extending the domain of featured-based model of property induction. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, IN.
- Smith, E. E. (1990). Categorization. In D. Osherson and E.E. Smith (Eds.), *An invitation to cognitive science vol. 3: Thinking*. (pp. 33-53). London: MIT Press.
- Smith, E. E., Shafir E., & Osherson D. (1993). Similarity, plausibility, and judgments of probability. In P. N. Johnson-Laird & E. Shafir (Eds.), *Reasoning and decision making* (pp. 67-96). Cambridge MA and Oxford UK: Blackwell.
- Smith, L. B., & Samuelson, L. K. (1997). Perceiving and remembering: category stability, variability, and development. . In K. Lamberts and D. Shanks (Eds.), *Knowledge, Concepts, and Categories* (pp. 161-195). London: Psychology Press.
- Spalding, T. L., & Ross, B. H. (1994). Comparison-based learning: Effects of comparing instances during category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 1251-1263.
- Springer, K. (1992). Children's awareness of the biological implications of kinship. *Child Development*, 63, 950-959.
- Schwartz, S. P. (1979). Natural kind terms. *Cognition*, 7, 301-315.

- Thagard, P. R. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12, 435-502.
- Tversky, A., & Kahneman, D. (1983). Extensional vs. intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293-315.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113, 169-193.
- Ullman, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.
- Wattenmaker, W. D. (1995). Knowledge structures and linear separability: integrating information in object and social categorization. *Cognitive Psychology*, 28, 274-328.
- Williams, E. J. (1959). The comparison of regression variables. *Journal of the Royal Statistical Society (Series B)*, 21, 396-399.
- Wisniewski, E. J. (1995). Prior knowledge and functionally relevant features in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 449-468.
- Wisniewski, E. J., & Medin, D. L. (1991). Harpoons and long sticks: the interaction of theory and similarity in rule induction. In *Concept Formation: Knowledge and experience in unsupervised learning*, D. H. Fisher, M. J. Pazzani, P. Langley (eds), 237-278. San Mateo, CA: Morgan Kaufmann.
- Wisniewski, E. J., & Medin, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, 18, 221-282.
- Wu, M., and Gentner, D. (1998). Structure in category based induction. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, Madison, WI, 1154-1158.

