Integrating the Mind

There are currently several debates taking place simultaneously in various fields of psychology that address the same fundamental issue: to what extent are the processes and resources that underlie higher cognition domain general versus domain specific? Extreme domain specificity argues that people are effective thinkers only in contexts which they have directly experienced, or in which evolution has equipped them with effective solutions. The role of general cognitive abilities is ignored, or denied altogether.

This book evaluates the evidence and arguments put forward in support of domain specific cognition, at the expense of domain generality. The contributions reflect a range of expertise, and present research into logical reasoning, problem solving, judgement and decision making, cognitive development, and intelligence. The contributors suggest that domain general processes are essential, and that domain specific processes cannot function without them. Rather than continuing to divide the mind's function into ever more specific units, this book argues that psychologists should look for greater integration and for people's general cognitive skills to be viewed as an integral part of their lives.

Integrating the Mind will be valuable reading for students and researchers in psychology interested in the fields of cognition, cognitive development, intelligence and skilled behaviour.

Maxwell J. Roberts is a psychology lecturer at the University of Essex.

5 Ontological commitments and domain specific categorisation

Steven Sloman, Tania Lombrozo, and Barbara Malt

Categorisation research is Janus-faced, with two orientations that rarely look at one another. One side is busy developing and testing algorithms for human classification that are presumed to be domain general (Ashby, 1992; Kruschke, 1992; Nosofsky, 1992), paying scant attention to the possibility that human categorisation processes differ depending on the kind of object being considered. The other side spends its time carefully documenting how categorisation differs across different object domains (e.g., Atran, 1998; Gelman, 2003), with little notice of the many well-specified and rigorous domain general models of categorisation that have been proposed.

One difference between the purveyors of these two literatures is their sense of what constitutes psychological explanation. The first group places more value on the precision and clarity afforded by formal models, the second on generalisability beyond the laboratory. But the contrasting assumptions of domain generality and domain specificity may reflect a more profound theoretical divide. Cognitive scientists have documented a variety of ways in which patterns of categorisation and induction differ as a function of domain. For example, children and adults tend to privilege external appearance and functional properties when reasoning about artifacts like chairs, but care more about internal properties and an entity's origin when reasoning about natural kinds like dogs (Keil, 1989). Differential patterns of categorisation and inference can be explained by appeal to domain specific representations or processes, and indeed, they frequently are. However, it is also possible that children and adults make differential judgments on the basis of properties that merely correlate with domain, like similarity (e.g., animals tend to have more in common with one another than with artifacts) or causal role (e.g., artifacts are created to serve human needs and natural kinds typically are not), without engaging representations or processes tailored to the different domains. If so, domain differences may tell us little about the most central aspects of categorisation.

In this chapter, we attempt to bridge the gap between these two orientations. We begin by spelling out the range of current theoretical views on domain specificity in categorisation, laying out claims about the origins of domain distinctions and the different mechanisms involved. The presentation will make explicit some of the assumptions latent in the various positions and allow us to evaluate the extent to which each finds support in existing research. To help specify the range of views, we offer a hierarchy of positions, beginning with the most strongly domain specific and ending with the most domain general. We take as our guide an analysis by Sloman and Rips (1998) of the role of similarity among objects and categories in human thought which distinguished four views: the extreme position that similarity is real, primitive, and cognitively efficacious in human reasoning (strong similarity), two variants of weaker claims, and the opposite extreme view that similarity is "invidious, insidious, a pretender, a quack" (Goodman, 1955), called "no similarity." We will suggest here that a set of positions on domain specificity in categorisation can be identified that also makes increasingly weaker claims, this time about the role of domain specific representations and processes.

At present, a compelling argument for invoking domain specific representations or processes in theories of categorisation comes from claims about psychological essentialism: that some categories are ascribed hidden, underlying essential properties that are causally responsible for an object's observable properties, and that these ascriptions play a central role in category judgments and inductions (e.g., Medin & Ortony, 1989). Most advocates of essentialism claim that only some categories are ascribed essences, and that such "selective essentialism" respects domain boundaries. However, debates about domain specificity turn on questions beyond essentialism. We have identified five key questions whose answers distinguish positions on domain specificity in categorisation (see also Table 5.1).

1. Are there innate modules that are differentially sensitive to objects from different domains?
2. Are there cognitive mechanisms (including representations, processes, or both) arising from any source that are differentially sensitive to objects from different domains?
3. Do people ascribe essential properties to (some) categories?
4. Are objects' causal roles critical to categorisation?
5. Are some domain general representations or processes operative in categorisation?

The strongest position we consider, *extra-strong ontology*, adopts the view that categorisation is governed by domain specific, innate modules that yield ontological kinds that are either essentialised or not. The weakest position we consider, *no ontology*, denies both domain specificity and essentialism. We also consider three intermediate positions based on claims in the literature, which leads us to the following five positions on the role of ontology in categorisation.

1. *Extra-strong ontology*. Domain differences in categorisation and induction are real and cognitively primitive (irreducible to other cognitive mechanisms). They result from innate modules evolved to pick out

Table 5.1 Properties of categorisation theories

	Determinants of category judgments and inductions				
Theory	*Innate modules*	*Domain specific mechanisms*	*Essentialist beliefs*	*Causal roles*	*Domain general properties*
Extra-strong ontology	✔	✔	✔	✔	✔
Strong ontology		✔	✔	✔	✔
Medium ontology			✔	✔	✔
Mild ontology				✔	✔
No ontology					✔

systematic regularities in the environment, including causal regularities. Modules have a one-to-one correspondence with domains (e.g., natural kinds versus artifacts), and have an all-or-none policy on essentialism: either all entities corresponding to a module are essentialised or none are. In a given classification or inference, an object is processed the way it is in virtue of the module that covers it.

2 *Strong ontology*. Domain differences are real and cognitively primitive. They result from learning mechanisms sensitive to properties like causal structure that partition objects in the world along deep ontological lines. Domains have an all-or-none policy on essentialism: either all entities from a domain are essentialised or none are. In a given classification or inference, an object is processed the way it is in virtue of the domain that covers it.

3 *Medium ontology*. Domain differences in categorisation and inference are real, but not cognitively primitive. They emerge from domain general, causal learning mechanisms that ascribe essences to some entities but not others. To the extent that the causal assumptions underlying the ascription of an essence correspond to causal differences across domains, essentialism will obey domain boundaries. In a given classification or inference, an object is processed the way it is in virtue of whether it is essentialised, which will correlate imperfectly with its domain.

4 *Mild ontology*. Domain differences in categorisation and inference are systematic, but not cognitively primitive. People reason and classify using domain general causal reasoning mechanisms. To the extent that domains correspond to causal discontinuities in the world, systematic differences between domains may emerge, and domain thus serves as a useful shorthand for theorists to classify roughly different types of processing. However, in a given classification or inference, an object is processed the way it is in virtue of its causal history and other causal roles, which will correlate imperfectly with its domain.

5 *No ontology*. Systematic domain differences in categorisation and inference may exist, but are not cognitively primitive nor due to causal

differences across domains. People classify and reason using domain general reasoning mechanisms. To the extent that domains differ with respect to the properties such mechanisms track, systematic differences between domains may emerge. However, in a given classification or inference, an object is processed the way it is in virtue of its non-causal, domain general properties, which will correlate imperfectly (if at all) with its domain.

These positions are ordered according to the strength of the claims they make: Stronger claims entail weaker claims, not vice versa (see Table 5.1). For example, according to extra-strong ontology, innate modules correspond to domains, which in turn determine whether an essence is ascribed. The modules serve the adaptive function of reasoning about a particular domain, and thereby encode systematic properties of the domain (Fodor, 1983, discusses what a module is). Such systematic properties might include the causal roles that entities in the domain typically play, as well as regularities that are not entirely causal, like the physical form of objects. Nonetheless, an advocate for extra-strong ontology will differ from other theorists in predictions about classification and inference in that judgments must be grounded in the relevant innate module, and not in essence or causal role.

While the theories in Table 5.1 entail the claims of theories lower on the hierarchy, myriad positions that do not obey this nested structure are logically possible. If each column is treated as independent, there are 2^5 positions, yet we only find five, with close variants, represented in the literature. Partly this is because the criteria in some columns have implications for others; partly it is a fact about theorists' predilections. Nonetheless, it is worth acknowledging the possibility of uncharted yet coherent views. For example, a theorist might advocate essentialism while denying the existence of meaningful domains, or advocate domain differences without appeal to essentialism. In fact, we have some sympathy with this latter position.

We next identify some of the theories associated with each of the five views and review the critical data that support or contradict each one. For each, we focus on evidence for and against the claim that distinguishes that position from its neighbour one step lower on the hierarchy. This review is not intended to be exhaustive but merely to clarify the claims and what they entail. Although this hierarchy can be fruitfully applied to a range of domains, in the interest of space we focus on the distinction between natural kinds and artifacts. We conclude by advocating the weakest position that can account for existing data: mild ontology supplemented with domain-level generalisations.

Extra-strong ontology

Extra-strong ontology is committed to the existence of modules that pick out domains and determine what properties found in those domains are essential. These modules are proposed to arise in humans through natural selection, because sensitivity to causal and other regularities in the environment provides a selective advantage. Extra-strong ontology is differentiated from strong ontology by the assumption that ontological distinctions result from *innate* mechanisms; that is, a module encodes some key structural aspect of categories that emerges regardless of exposure to objects in the relevant domain. While people need not represent ontological kinds explicitly (although they might), the view requires that different ontological kinds (primarily living things) be treated in a unique way by the processes that assign category membership and make inductions. People might even encode differences that do not correspond to actual differences in the world, such as race (Templeton, 1998) and biological species (Mayr, 1982; Sober, 1994).

This view is articulated most fully by Atran (1998), who claims that a distinct module exists to make classifications and inductions related to plants and animals, and that this evolved via natural selection to reflect the history of interaction between people and these sorts of ontological kinds. He writes: "Universal taxonomy is a core module, that is, an innately determined cognitive structure that embodies the naturally selected ontological commitments of human beings and provides a domain-specific mode of causally construing the phenomena in its domain" (p. 555). Atran claims that people automatically generalise biological claims about an object to the generic-species level, a level that corresponds both to biological genus and species. He proposes that people make the "commonsense assumption that each generic species has an underlying causal nature, or essence, that is uniquely responsible for the typical appearance, behaviour, and ecological preferences of the kind" (Atran, 1998, p. 548), and that this assumption is innate. By "essence," he means "an intrinsic (i.e., nonartifactual) teleological agent, which physically (i.e., nonintentionally) causes the biologically relevant parts and properties of a generic species to function and cohere 'for the sake of' the generic species itself" (pp. 550–551). Note that essences need not be specific, identifiable parts or attributes; the claim is that people believe *something* essential exists whether or not it actually does and whether or not they have more specific beliefs about it. This position is echoed by Pinker (1997).

Support for extra-strong ontology

Categories in different domains do tend to have different causal histories. Living things evolve within ecosystems whereas artifacts are the product of intentional human design. To the extent that natural selection shapes the mind to appreciate such aspects of the world, natural selection might build in this distinction. It would be particularly important if the information needed to be accessed quickly and automatically.

Beyond such speculation, the assumption of hardwired domain specific modules is consistent with the following observations.

1 Living things are categorised in largely the same way cross-culturally (Atran, 1990; Berlin, Breedlove, & Raven, 1973; Malt, 1995). Atran concludes there is a universal "general-purpose" taxonomy of living things (though see Ghiselin, 1998). Universality would be a natural, though not logically necessary, consequence of innate modularity.
2 Some evidence indicates a cross-cultural tendency to prefer inductive inferences that are made at the generic-species level (e.g., vulture), even when this does not correspond to the basic level of categorisation as determined by other measures, like naming and feature listing (Coley, Medin, & Atran, 1997). For example, people are almost as willing to project a blank biological predicate (like "has enzyme X") to the generic-species level as they are to more specific levels, but they hesitate to project such properties beyond the generic-species level. This finding holds true not only for populations like the Itzaj Maya, for whom the generic-species level is the basic level of categorisation, but also for American college students, who typically treat the life-form level (e.g., bird) as the basic level on tasks other than inductive inference, like naming. This inductive tendency can be explained by appealing to an innate tendency to generalise at the generic-species level.
3 Evidence has been reported for selective neuropsychological impairment of knowledge about animals (e.g., Caramazza & Shelton, 1998), consistent with the claim of an innate module that must have some form of neural representation.
4 Even young children show systematic differences when making inferences about natural kinds versus artifacts. In particular, nonobservable internal features are given more weight in classification and induction of living things and observable external features are given more weight for artifacts (reviewed in Gelman, 2003).
5 The animate/inanimate distinction is made very early in development, in infancy (Rakison & Poulin-Dubois, 2001).

Problems with extra-strong ontology

Although a lot of data are consistent with extra-strong ontology, this view makes strong assumptions while providing little more than a description of the data, rather than an explanation in terms of cognitive mechanisms. The notion of module merely mirrors the observation of domain differences without explaining how they arise or offering novel predictions. Any unambiguous support for the extra-strong ontology view must come from some other source, perhaps an evolutionary argument for innateness as offered by Atran (1998). Evaluating such an argument is beyond the scope of this chapter. When doing so, of course, it is critical to consider not just whether a plausible evolutionary account can be given of the modular view, but also whether one can be given of its antithesis, a non-modular view (see Sloman & Over, 2003; Sterelny, 2003).

Extra-strong ontology also has difficulty accounting for cross-cultural variation in biological kind classification and inference. Despite broad similarities, animal categorisation exhibits some cultural specificity: For instance, cows have a religious significance that sets them apart in India, whereas they are often grouped with sheep and goats and other livestock in Wisconsin. The flexibility of biological categorisation has led theorists to propose promiscuous realism (Dupré, 1981), the idea that there is no single, privileged way to carve up kinds in the environment. Variation in animal categorisation is not arbitrary – it tracks the different roles that animals play in the economic, social, and religious lives of different cultures (Diesendruck, Markson, & Bloom, 2003, and Malt, 1995, provide references). Extra-strong ontology could account for cultural variation by assuming the module that governs biological classification is sensitive to complex causal roles that crosscut domains. But to the extent that categorisation cuts across domains by virtue of learning these causal roles, the innate and domain specific components of the module cease to account for the data.

Moreover, taxonomies do not always support induction. Medin, Coley, & Storms (2003) provide several examples of inductions mediated by causal reasoning that break category boundaries. Sloman (1998) reports cases where people fail to use category hierarchies to make inductions in the most transparent situations. For instance, people do not necessarily infer that all sparrows have a property when told that all birds do, even when they affirm that sparrows are birds. The problem for extra-strong ontologists is that the only inferential mechanism they propose is taxonomic, and yet taxonomic reasoning is surprisingly rare.

Finally, views like extra-strong ontology provide no natural account of expertise that transcends evolutionary endowment. For example, biologists appeal to categories and processes that reflect non-obvious scientific discoveries, and have even redefined domain boundaries (e.g., by interpreting life as a biological rather than a spiritual phenomenon). If such knowledge emerges from the same mechanisms that govern folk biological inference, then the constraints imposed by mechanisms in the biology module must be too weak to warrant strong claims about the domain specific nature of categories and inference. Alternatively, such knowledge could result from inferential mechanisms outside the folk-biology module. But if such non-modular inferential mechanisms are available, why have a module for folk biology in the first place?

One possibility is that innate, domain specific representations and processes are needed to get folk biology off the ground, but that once an initial scaffold of biological knowledge is in place, other mechanisms take over and account for scientific expertise and cultural variation. Proposals that such representations are required at the outset are often based on the observation that perceptual information is insufficient to explain people's categories. For example, children might be observed not to treat toy dogs and real dogs in the same way; therefore theorists conclude that they use domain-dependent conceptual knowledge to make critical distinctions. This argument is based on treating domain specific conceptual processes and domain general perceptual heuristics as the only possible categorisation processes (Atran, 1998, makes this leap explicitly on p. 554). Clearly, the dichotomy is not exhaustive. In our discussion of mild ontology below, we offer an alternative in terms of domain general conceptual processing.

Strong ontology

The idea of strong ontology is that, as a result of an essentialist bias and causal learning that picks out regularities in the causal structure of the world, people have domain specific mechanisms for classification and inference. Strong ontology differs from extra-strong ontology in denying innate modules, but shares the assumptions that domains are associated with an all-or-none policy on essentialism and that the entities from some domains (like living things) are essentialised while others (like artifacts) are not. As Gelman, Coley, and Gottfried (1994) put it, essentialism "readily applies to new domains that have never before been encountered" but does not "sensibly apply to all domains . . . An essentialist assumption may be 'designed' in such a way that it functions only when it meets domains with the appropriate features" (p. 358). Strong ontology differs from medium ontology in the assumption that ontological distinctions emerge from domain differences, and that essentialism respects these differences.

At various points, strong ontology seems to be the view espoused by Gelman (2003). She writes: "Essentialism is a domain-general assumption that is invoked differently in different domains depending on the causal structure of each domain" (p. 312). However, the result of differentially applying these domain general assumptions seems to be domain specific, essentialist reasoning: "Although the proposed principles are domain-general, essentialising is not. We do not essentialise wastebaskets or gum-balls" (p. 321). At other points, Gelman (2003) seems more consistent with the position we review in the next section, medium ontology.

Support for strong ontology

The main advantage of this approach is that it too is consistent with all of the phenomena used to motivate the extra-strong ontology view, but is unburdened by assumptions of innateness or those concerning modularity. The supposition that some domains and not others are "essentialised" is intended to account for the very data that motivate extra-strong ontology. And by allowing a role for learning, strong ontology is consistent with the cross-cultural variation in biological classification that threatens extra-strong ontology.

Problems with strong ontology

Strong ontology has difficulty accounting for cases in which classification and induction do not respect domain differences. After all, objects – including living things – play many causal roles that are not tied directly to their biological domain. Hence inferences often depend on contingent aspects of the object, its environment, and the specific goal of the inference. For instance, a dog plays many causal roles in the day-to-day activity of a family, including drawing family members outside and causing them to administer medication. Some inductions, for example, about behaviour and disease, will depend in part on these other causal roles. Only a few, like those having to do with body parts, will be based purely on the dog's biological history (and even then, not if she loses a leg in a car accident). A dog's disposition may be as much a result of a boisterous household as of her evolutionary niche. Moreover, some important groupings have nothing to do with ontological domain. The class "things to take out of a burning house" (Barsalou, 1991) includes living things (like a dog) and artifacts (photos and jewellery), but not cockroaches and wastebaskets. To account for such inferences and groupings, strong ontology is forced to appeal to ad hoc alternative explanations.

Moreover, essentialist assumptions do not cleanly line up with domain assumptions. For instance, it is natural to distinguish the domains of living things and artifacts, yet one might be prone to assume that both living things and some artifacts have essences. This tension is clear in Keil's work (Keil, 1995). His position is close to strong ontology for he asserts a set of dimensions that normally distinguish natural kinds and artifacts, but he allows exceptions: "[T]hese contrasts with natural kinds turn out to be more like rough rules of thumb than strict criteria. Artifacts and natural kinds appear to be arrayed along several related continua rather than in sharply contrasting bins, as is seen with more complex artifacts, such as televisions, cars, and computers, and with designed living kinds, such as plants and animals subject to intensive breeding" (Keil, 1995, p. 235). Indeed, some theorists assert that all artifacts have essences, namely the intent of the artifact's creator (Bloom, 1996; Matan & Carey, 2001). Of course, if artifacts do have essences, the value of essentialism for explaining domain specific phenomena that distinguish artifacts and living things is undermined. Essentialism could account for such phenomena by positing that the essences of artifacts and living things have different content, but then it is the content, and not essentialism, that does the necessary explanatory work.

The one weakness of the strong ontology programme relative to extra-strong ontology is the inability to explain the very early acquisition of domain specific beliefs because strong ontology entails a period of causal learning of domain differences (see Rakison & Poulin-Dubois, 2001). Quinn and Eimas (1996) found that young infants distinguish animals and artifacts. Pauen (2002) has shown that 8-month-olds distinguish animals and furniture, even when some perceptual properties are equated across the two domains. However, these results might reflect perception of differences across those dimensions that were not equated (like texture), and not necessarily domain specific knowledge.

Medium ontology

Medium ontology allows that, although people are assumed to essentialise, essentialist beliefs do not necessarily line up neatly with domains. Inferences are made by virtue of generalisations over domains only insofar as domains pick out causal regularities, not by virtue of the domains per se. We believe that medium ontology represents the modal view in the literature on cognitive development, as many authors assume that a belief in essence is critical to categorisation processes while denying that essences are invariably assumed in some domains but not others (Keil, 1989, 1995; Kelemen & Carey, 2007). This is the view found, for example, in Gelman and Hirschfeld (1999). They argue that the "early and nearly parallel emergence of essentialist reasoning in these different domains is consistent with the maturation of a single conceptual bias for essentialist reasoning" (p. 421), but they suggest that the essentialist bias crosses domains: "Essentialism may not map cleanly onto domains. Events and specific entities . . . may be essentialised without essentialising the larger domain of which it is part" (p. 437).

Essentialist beliefs are described in various ways by various authors (e.g., Keil, 1989, 1995; Medin & Ortony, 1989). Gelman and Hirschfeld (1999) say a causal essence is "the substance, power, quality, process, relationships, or entity that *causes* other category-typical properties to emerge and be sustained and confers identity" (p. 406, emphasis in original). Gelman (2003) identifies essentialism as a set of assumptions that people, including young children, make about various domains, as follows.

1. Appearance is distinct from reality. This explains why non-obvious properties of objects can be the most central.
2. Properties come in clusters that support induction.
3. Properties have deterministic causes. The most central properties are root causes.
4. Objects maintain their identity over time. Therefore, understanding an object's origin is important.
5. People must defer to experts because one does not necessarily know an object's causal properties or origins. An essence is merely a placeholder for an ultimate cause. This forces people to accept some category anomalies.

Other construals of psychological essentialism bear a family resemblance to Gelman's definition. The notion of psychological essentialism is related to but distinct from Putnam's (1975) and Kripke's (1980) notions of linguistic essentialism.

Rips (2001) points out that an essence can only serve to distinguish categories on the assumption that it constitutes a necessary and sufficient condition for category membership. Gelman (2003) calls such a defining set of features a "sortal essence" and denies that the causal essences she attributes to human thought have this property. But if a causal essence were not necessary for category membership, then there would be category members without the causal essence that would have few if any properties of the category because they would be missing the root cause of those properties. And if the causal essence were not sufficient for category membership, then there would be nonmembers of the category that nevertheless have the category essence. Because they have the causal essence, they would have most of the properties of the category and yet not be category members. Neither of these conclusions makes sense. Therefore, denying necessity and sufficiency seems to undermine the work that essences purport to accomplish, and the property of being defining should be added to Gelman's list.

Support for medium ontology

This position offers greater explanatory depth than the previous positions. It recommends a theoretical programme: that of spelling out the dimensions on which categories differ and specifying precisely how those dimensions influence classification and induction decisions, and when and how essentialist beliefs play a part. Notice that the explanatory work is done by the assertion that causal regularities are learned and not by the assumption of essentialism. The value of the essentialist assumption may be to explain what motivates people to learn those regularities.

Problems with medium ontology

Like the innate modules of extra-strong ontology, the essentialist claim seems more of a restatement of a set of phenomena than an explanation. Calling Gelman's (2003) list of insights about human thought "essentialism" does no more than label the insights. The absence of a coherent theory that puts the phenomena together in a principled way may be why different theorists draw essentialist lines in different ways.

Furthermore, several experiments have tried to operationalise the notion of essence and found it wanting. For instance, Malt (1994) found that people's beliefs about the presence or amount of H_2O in liquids did not serve to distinguish what they did and did not name water. Kalish (1995) found that people were not willing to treat even natural kinds as having absolute category membership. Braisby, Franks, and Hampton (1996) found that the presence or absence of normal cat essence did not determine whether people classified as cats animals that, in some cases, turned out to be robots controlled from Mars. Gelman and Hirschfeld (1999) argue that all these operationalisations are inadequate in one way or another, but this raises the question of the value of a concept that is so difficult to operationalise.

One major weakness of this view, along with its stronger counterparts, is its assumption that essentialist beliefs are at least sometimes operative in the categorisation process. Several authors have argued that causal knowledge can do all the explanatory work done by essentialism without any appeal to essences per se (Rips, 2001; Sloman & Malt, 2003; Strevens, 2001). Arguments in favour of medium ontology often rely on the same dichotomy between domain specific conceptual processes versus domain general perceptual heuristics that one finds among extra-strong ontologists. One argument has the form "People don't classify or induce based on perceptual similarity, so they must be relying on a hidden essence. And essences have a different status in different domains, therefore classification and induction are domain specific." But this argument ignores the mild possibility, discussed below, of domain general conceptual processes that result in different patterns of categorisation and reasoning across domains.

Mild ontology

This view assumes neither that essentialist beliefs guide our classifications and inductions, nor that classifications and inductions are done differently in different domains. Instead, people have reasoning mechanisms that use whatever causal knowledge is available as a means of explaining how the world works. As in medium ontology, the assumption is not that objects are processed by virtue of their ontological domain, but rather that objects from similar domains tend to be processed in the same way because causal regularities happen to be associated with domains in the world.

Instead of positing essentialism, this programme simply assumes that people are interested in explaining the world around them and they build causal models to do so. These are assumptions that must be made by medium and strong ontology as well. A more detailed proposal about how people use causal regularities to classify and make inductions follows. Its assumptions require neither essentialism nor domain specificity and yet are consistent with the phenomena described so far. They are motivated by the literature on causal Bayesianism; Sloman (2005) offers an introduction and Gopnik and Schulz (in press) review recent advances. Scholl (2005) offers a related Bayesian (though not causal) resolution to the innateness/learning debate.

- *Causal determination*. Events and properties are assumed to have causes.
- *Inductive leverage*. Causes and effects relate to other causes and effects in systematic ways that are measurable. For one, certain patterns of causal relations map onto patterns of dependence and independence in probability distributions. For example, the heart's pumping causes oxygenated blood to travel to the limbs, which in turn causes transfer of oxygen to cells in the limbs. This explains why there is a statistical dependence between the heart pumping and oxygenated cells in the limbs. It also predicts that knowing whether or not blood is travelling to the limbs would render the other two independent. As a result of this mapping between causal structure and statistical dependence, causal structure can be induced from observation and experimental intervention. The structural properties of causal relations include such things as "explaining away": The presence of one explanation for an effect decreases the credibility of a second, independent explanation.
- *Multiple scales*. People are sensitive to causal structure at multiple scales. For instance, they learn and use structure involving coarse generalities like "evolution causes adaptations" as well as structure involving more specific mechanisms like "molars are for chewing."
- *Precision/error trade-off*. To make inductions, people choose a level of causal structure that maximises inductive strength. This involves trading-off precision and accuracy. Coarse causal knowledge is likely to lead to correct but imprecise inference (e.g., organs that evolve through natural selection are likely to have *some* adaptive function). Fine-grained causal knowledge will lead to precise conclusions, but is more prone to error (e.g., if dogs use molars for chewing, then all animals use molars for chewing). In the language of statistics, the proportion of variance accounted for must be traded off with the number of parameters required by the model. Specific models may capture more variance, but they make more assumptions and therefore are more likely to generalise incorrectly. More generally, the best theory maximises empirical coverage (and is therefore specific) while maintaining explanatory simplicity (and is therefore not too specific).
- *Pragmatic tailoring*. Individuals (aided by culture) are pretty good at maximising inductive strength. For different tasks, this maximisation leads to different levels and, more generally, different relevant causal structures. The relevant structure for naming might involve satisfying some functional goal (I want to open the door, so the function of the handle is important to selecting a name) or might involve reference (I want you to hand me that particular object and so I name according to physical properties). The causal structure for making inductions about behaviour could have to do with biological predispositions (Labradors like to retrieve) or with environmental context (dogs that belong to that surly guy are all vicious).

Mild ontology interprets causal relations broadly. Any causal role can potentially be relevant to classification and induction. A view of this type is offered by Barsalou, Sloman, and Chaigneau (2005), who use a graphical probabilistic framework to describe the HIPE theory of the multiple levels of causal knowledge that people have about various categories. The theory describes what people know at a high level of abstraction, knowledge that cuts across domains. People believe that every object has a history (H) that causes its physical make-up (P) that, along with other necessary causes, also causes functional events (E): The "I" in HIPE reflects the fact that this knowledge reflects an intentional perspective.

$$H \rightarrow P \rightarrow E$$

This generic domain general knowledge implies that nonobservable features relating to the past play a central role in the category's structure, suggesting that domain-dependent knowledge is not required to explain the importance that both children and adults put on nonobservable properties when categorising. Moreover, no property plays any essential role in this causal model. What makes a representation viable is the coherence of its causal relations. Some support for the relevance of domain general reasoning comes from Schulz and Gopnik (2004), who showed that 4-year-olds can make inferences about causal structure from probabilistic data, and that the inferences are domain general in that they are made in multiple domains and across domains.

Knowledge also exists at a lower level of abstraction according to HIPE; H, P, and E can all be more fully specified for particular domains, like living things and artifacts. A summary form of slightly more detailed models is shown in Figure 5.1.

Hence, according to the theory, whereas different domains are associated with different models at a particular level of abstraction, they are associated with the same model at higher levels. More important, those models are just rough characterisations of the relations among relevant causal mechanisms. The models are described using a generic language in which nodes represents events and properties, and links represent causal mechanisms. Different models could be relevant depending on the reasoner's purpose; sometimes they might be unrelated to the models depicted above.

Rehder (in press) offers a more detailed model that appears consistent with mild ontology, though at a single scale. Although he is willing to talk about "essences," Rehder's idea is that classification is mediated by knowledge of generative causal structure and that properties are important to the degree that they are consistent with causal knowledge. Rehder and Hastie (2004) offer a model of induction that appeals to causal structure without any direct appeal to belief in essences or to domain specific knowledge.

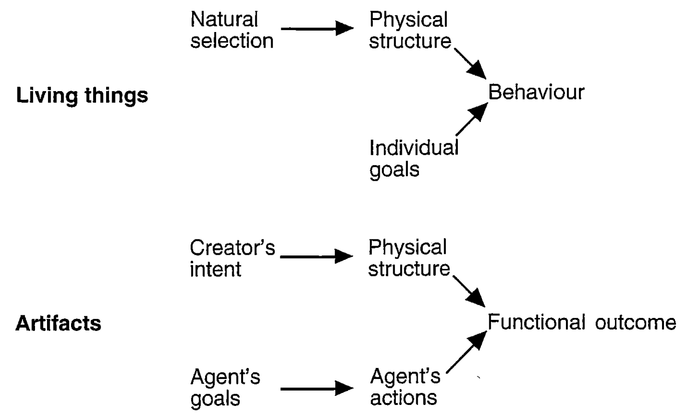

Figure 5.1 Examples of HIPE models with relevant details specified.

Support for mild ontology

This view inherits the advantages of medium ontology while making even fewer assumptions. Medium ontology assumes that people have an essentialist bias and also that they are biased to learn causal regularities in the world. Mild ontology accepts the centrality of causal structure and differs from medium ontology merely in denying that this structure arises from essentialism. Because medium ontology's explanations for the phenomena are all based on causal beliefs, not on essentialism per se, mild ontology inherits all medium ontology's explanations for the data so far discussed. That is, it assumes that children quickly learn causal regularities and use them to classify and make inductions. The only difference is that medium ontology says that children do it due to an essentialist bias, whereas mild ontology says they do it by virtue of a desire for causal explanation. As medium ontology likewise needs to assume that children seek explanations in order to pick out causal regularities, mild ontology seems to be making one fewer assumption.

To see the power of the causal learning assumptions of mild ontology, consider how they explain preference for the generic-species level for inductive projection: That level is a frequent (though not exclusive) level of preferred projection for living things due to the convergence of causal structure at that level. The convergence is largely due to the fact that this is the level that supports procreation, which has enormous causal significance with regard to evolution, common anatomical and physiological structure, genetic disease, behavioural adaptation, as well as determining current goals, social behaviours, family structure, etc. An enormous set of causal properties revolves around procreation, and that is what people are responding to when they indicate a preference to project across the generic-species level, not beliefs about the inductive potential of the level of generic species per se.

Mild ontology has the additional advantage that it is consistent with the deviations from Atran's (1998) folk-biological taxonomy. For instance, Coley et al. (1997) found that people's preferred level for naming could be higher than their preferred level of induction, the generic-species level. A causal constraint on naming that arises from its communicative demands is that all interlocutors must know the term being used. This is not a constraint on induction. Hence, the basic level of naming will sometimes be at a higher level than induction owing to differential causal constraints imposed by different categorisation tasks. A more complete analysis of the various tasks that define the basic level would reveal other differences. For example, the feature listing task (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) depends on knowledge that can be articulated whereas induction does not (people might know not to generalise from oaks to elms even without being able to specify how they differ). For a fuller analysis of categorisation task constraints, see Malt and Sloman (2007).

The relatively weak assumptions of mild ontology are also necessary to explain the fact that induction is often mediated by causal relations that have nothing to do with ontological knowledge. For example, Heit and Rubinstein (1994) show that induction depends not just on ontological kind, but on the requirements to support particular predicates (e.g., animals that eat at night require certain sensory capacities regardless of their species). Medin et al. (2003) found that people make causal inferences using variables that have no correspondence to folk-biological taxonomies. Inductions can be based on nontaxonomic relations like containment (diseases are more likely to be transferred from mice to owls/cats than from mice to squirrels, because owls and cats eat mice).

Mild ontology is also consistent with the many different ways that entities in any domain can be grouped into categories. For instance, in the domain of natural kinds, tomatoes and mushrooms might be categorised together as sauce ingredients by virtue of their causal properties with respect to human taste. Similarly, Malt, Sloman, Gennari, Shi, and Wang (1999) showed that artifact naming is highly language-specific, a fact that has to do with different linguistic conventions in different cultures. The different naming patterns might reflect – at least in part – different causal histories in different societies. The fact that French has no single word that corresponds to the English "jacket" may reflect different societal differences in attitudes toward function and fashion. It also turns out that the judged causal centrality of properties can be manipulated by introducing new causal relations among properties both in classification (Ahn, 1998) and in induction (Hadjichristidis, Sloman, Stevenson, & Over, 2004): A property is judged more central when judges are told that other properties depend on it.

These effects imply a flexible use of causal structure in classification and inference. To be fair, essentialism does not rule out the possibility that objects can be grouped in multiple ways. But it does imply that one grouping has some priority, namely, the grouping that conforms to the proposed essences. If other, non-essentialist groupings also exist, there must be cognitive machinery to create and use these. We have claimed that this cognitive machinery involves causal models and that it can explain both "essentialist" and "non-essentialist" groupings. At this point, the onus is on the essentialist theorist to explain why additional machinery is required for the "essentialist" cases.

Problems with mild ontology

Such broad construal of causal roles requires the mild ontologist to be explicit about causal representation and learning, because the causal roles that properties can play are too varied and unpredictable to have been encoded directly by evolution. So work must be done to explain how causal structure is learned and deployed on line. Of course, a cognitive account of any of the phenomena requires this work anyway, because an appeal to evolution does not provide any direct understanding of cognitive processes or representations that govern categorisation.

Atran (1998) argues that the weakness of views like those encompassed by mild ontology (and probably medium ontology as well) is their inability to explain the "global relationships linking (e.g., generating) species and groups of species to and from one another" (p. 566). The point seems to be that cultures and individuals differ immensely in their degree of integration of causal principles and in their knowledge of causal structure, yet folk biology is largely universal. Indeed, much causal knowledge is extremely vague (Keil, 2003), making it hard to see how learning would result in common principles of taxonomic classification and inference.

However, causal knowledge exists at multiple scales and even vague knowledge is knowledge. Few of the details of procreation need to be understood to appreciate that anatomical, physiological, behavioural, etc. systematicities are likely to be shared at the level that generally supports reproduction, the generic-species level. Even someone who lacks understanding or who does not believe that the generic-species level supports procreation will note all the correlated structure at that level and will thus give it priority.

Perhaps a more critical problem with the mild ontology view is that it provides no account of why people have an essentialist intuition. That is, why do people assert that some objects have a true nature that determines their kind? Why do people have disputes about whether a mule is "really" a donkey or a horse?

The mild ontologist would first note that people's intuitions are hardly clear on this point. In a systematic study of people's assertions, Kalish (1995) concluded that people assert that membership in even animal categories is a matter of degree. But to the extent that intuitions about the existence of essences are real, they do not necessarily arise from the same source as most everyday categorisation decisions (Armstrong, Gleitman, & Gleitman, 1983).

Intuitions that categories have essences could arise as an overextension of the prerequisites of communication. Communication requires that assertions have truth conditions. Successful communication depends on interlocutors assuming that either a message passer believes their message to be true or the person has some rational motivation for passing it anyway. So conversation and other forms of communication take place under the usually implicit assumption that the categories under discussion are real and enduring and that they correspond to the way the world actually is. What people often neglect is that what determines whether an assertion is true or false is just a matter of convention with respect to a particular community. People select groupings for objects, and names to capture those groupings, in order to conform to whatever conventions are relevant at the moment (conventions that sometimes include rigorous scientific criteria). When probed, instead of ascribing the validity of an assertion to these conventions, for the sake of simplicity people ascribe it to some deeper, context-independent, essential property that may play no role either in the grouping or its name. For example, English speakers distinguish "fruit" from "vegetable" via a convention related to the role of the food in their meals (derived from properties like taste and texture), not from deeper botanical or biological properties (Malt, 1990). But people tend to believe that there is some deeper commonality that makes an apple and an orange appropriately called "fruit" in everyday English, a property not shared by eggplant or zucchini. A person's affirmation that a song is typical rhythm and blues stems from the belief that there is a set of people with expertise in music who would concur. The label indicates a desire to conform to a community. The essentialist bias refers to the persistent further intuition that the convention reflects some essential quality of R&B whether or not it does. It is simpler for people to believe that a category has a reality across contexts, and therefore an essence, than to specify the context-specific conventions that support a given assertion.

No ontology

The theories reviewed so far have not been formalised, yet a number of mathematical theories of the categorisation process have been proposed. With the exception of Rehder (in press), none makes any reference to essences or to domain specific processes (Ashby, 1992; Kruschke, 1992; Nosofsky, 1992). These theories do not make reference to ontological distinctions (for an exception see Lamberts & Shapiro, 2002), and the processing principles they advocate are completely domain general. Depending on the theory, these processing principles include the following: categories maximise the ability to predict new features (Anderson, 1991), category boundaries maximise separation among distributions of instances (Ashby), and categorisation involves comparison to exemplars (Nosofsky, Kruschke). Although they have not addressed the domain specificity issue, the implicit claim of all these models would seem to be that domain differences do not have systematic effects on the central processes of categorisation, or at least that no explanatory leverage on processes of classification and induction is gained by drawing distinctions among ontological domains.

Other theorists have proposed that classification and induction are mediated by a general causal reasoning system without claiming that the system is specifically tuned to domain differences (Lien & Cheng, 2000; Sloman, 2005). Of course, to the extent that such hypotheses would lead to models that learn causal structure that systematically varies across domains, they will effectively be implementing mild ontology.

Support for no ontology

The no ontology position inherits the advantages of mild ontology. By virtue of being completely domain general, it is consistent with the deviations in biological classification that favour mild ontology.

Problems with no ontology

Although the no ontology theories offered do not say anything inconsistent with the data reviewed, this position offers no leverage to explain the differences that do exist between natural kind and artifact categories. Many of the claims about categorisation made by these models have not been tested with objects from different domains, so their generality is unknown. In fact, most of the tests of these ideas have not even used real categories, but artificial ones based on simple physical attributes.

The data we have reviewed do seem to imply that a notion of causal structure is indispensable for explaining facts like those concerning the relative importance of internal and external properties. If differences in causal structure were uncorrelated with domain differences, then knowing an object's domain would give no purchase on understanding how people classify and make inductions with it, and the no ontology approach would be viable. But causal structure is correlated with domains, insofar as there are overall differences between the causal roles of living things (e.g., they are shaped by natural selection, they tend to be composed of homeostatic systems, see Keil, 1995) and artifacts (e.g., they are created by people, they tend to serve a human function). The mild ontology view recognises these differences and provides a means of talking about them.

The no ontology theorists have the equipment to become mild ontologists. They merely need to state explicitly how their theories distinguish the distributions of properties of, say, living things and artifacts. A superficial account could assert that there are distributional differences by stipulating, for instance, that nonobservable features are given relatively more weight in categorisation than observable features, for living things over artifacts. A richer explanation would provide an account of how that difference comes about. However, doing so would require an account of causal learning and representation, something that most no ontology theorists have not offered.

Conclusions

Because we see no motivation for essentialism separate from causal knowledge, we do not subscribe to medium, strong, or extra-strong ontology. We do believe, however, that there are important distributional differences in the way people represent living things, artifacts, and surely other domains, such as nominal kinds. Therefore we believe that a stronger stand than no ontology is called for. This leaves mild ontology. And in fact the best explanation for the evidence that we have covered is that people use specific causal knowledge to make inductions and to classify when they can.

However, there is reason to suppose that people use generic domain knowledge when lacking specific causal knowledge. For example, Goodman (1955) posits beliefs that mediate induction called "overhypotheses" that can hold over a fairly abstract domain. An example might be "different kinds of animals have characteristic mating behaviours." Such beliefs are clearly tied directly to knowledge about a domain. Shipley (1993) provides evidence for the psychological reality of such beliefs in the process of induction.

Further evidence that taxonomic knowledge serves as a fallback in the absence of specific causal knowledge can be found in comparing novice to expert induction. American undergraduates without much knowledge of trees show the diversity effect: Evidence is treated as warranting projection to the degree that the evidence comes from dissimilar sources (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990). For instance, when told that two types of tree are susceptible to a disease, the extent to which they project susceptibility to a third type depends on the degree that the first two types are dissimilar. Tree experts, in contrast, tend not to show the diversity effect (Coley et al., 1997). Instead, they reason in terms of ecological niche. When people have specific causal knowledge, they reason on the basis of it. But when they do not, they use taxonomic knowledge that may be tied either to a domain like living things (as Osherson et al. suggest) or to a feature-based representation that uses more generic knowledge (Sloman, 1993).

A final point in favour of mild ontology augmented by domain-level generalisations or overhypotheses comes from its generality. The proposal nicely accommodates the data we have reviewed on categorisation and induction, but also extends to other judgments, like the appropriateness of causal explanations. For example, most Western adults judge teleological explanations – explanations in terms of a function or goal – as appropriate for biological kinds, but not other natural kinds (Keil, 1992; Kelemen, 1999): one says that people have eyes "for seeing," but not that there are clouds "for raining." These judgments appear to be a function of the domain whenever tasks do not provide participants with a specific causal history for the property being explained. For example, Kelemen (1999) asked children and adults questions like "why are rocks pointy?," and provided a teleological and a non-teleological explanation. Adults systematically chose the non-teleological explanation for non-biological kinds like rocks, but the teleological explanation for artifacts and biological parts. This is presumably because they made an inference about the causal process likely to have led to what was being explained, and did so on the basis of the object's domain: a function-driven process like natural selection is a plausible cause for biological parts, while seemingly random physical processes are plausible causes for non-biological natural kinds. Lombrozo and Carey (2006) presented adults with similar questions, but for which the causal history of the property being explained *was* provided. In such cases, judgments of the appropriateness of a teleological explanation depended only on the causal history provided, with no effect of the object's domain. Such findings suggest that previously documented domain-level differences deserve closer scrutiny in terms of underlying causal differences.

The suggestion that people use kind membership when they do not have the specific causal knowledge needed to categorise is analogous to Quine's (1970) claim that similarity devolves into theory given enough knowledge. Quine argued that people construct categories based on similarity (e.g., whales go with fish) until they have enough theoretical knowledge to support their categories (many people come to believe that whales are mammals, not fish). We are suggesting that, analogously, ontology devolves into causal structure given enough knowledge. People use causal knowledge when they can; ontological knowledge serves as a proxy otherwise. The value of domain specific knowledge depends on what kind of causal structure could mediate the particular induction or classification at hand. If an induction involves what your mother would like her coat to be made of, or what your colleague would like to be served for dinner, domain specific knowledge at the life-form level is relevant (e.g., plants are OK, animals are bad) and maybe at the specific level (Spanish tomatoes are good, Northern Ontario tomatoes are bad), not necessarily only at the generic-species level.

In conclusion, we have tried to specify how domain specific knowledge is relevant to processes of categorisation, based on an analysis of what is required to explain the extant data, and not merely on intuition. We freely admit that the belief that certain kinds are what they are by virtue of an underlying essence can be intuitively compelling. Kalish (1998) showed that both children and adults treat both animals and artifacts as objective at the basic level, and they treat animals as more objective at the superordinate level. What we deny is that this intuition tells us very much about the cognitive processes associated with classification and induction. People can have content in their beliefs about category membership that is not utilised in their categorisation decisions. Malt (1994) has demonstrated this for a natural kind: Whether people call a substance "water" is independent of their beliefs about how much H_2O it contains. Malt et al. (1999) have demonstrated it for artifacts: The way people sort containers and dishes is not determined by the names they give the objects. Despite people's claims about the properties that constrain the categories they form, the actual categories are influenced by factors they are not aware of. As a result, the content of the beliefs and knowledge people hold may differ, and yet processes of categorisation and inference may be universal.

Hence we view the contrasting assumptions of the two traditions of categorisation research as a real theoretical divide but not an unbridgeable one. We see the hope of a rapprochement between those who favour formal theories focusing on domain general mechanisms and those who grapple with the messier real world and who have been primarily concerned with essentialism and domain specificity. Greater appreciation for the role and nature of causal reasoning would give the formal theorists the power to begin to explain the deep insights into the nature of human categorisation that have come from the other side.

Acknowledgments

The authors would like to thank Tom Griffiths and Daniel Weiskopf for helpful comments on an earlier version of the manuscript.

References

Ahn, W. K. (1998). Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition*, *69*, 135–178.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429.

Armstrong, S. L., Gleitman, L. R., & Gleitman, H. (1983). On what some concepts might not be. *Cognition*, *13*, 263–308.

Ashby, F. G. (1992). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449–483). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Atran, S. (1990). *Cognitive foundations of natural history: Towards an anthropology of science*. Cambridge: Cambridge University Press.

Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavioral and Brain Sciences*, *21*, 547–610.

Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 27, pp. 1–64). San Diego, CA: Academic Press.

Barsalou, L. W., Sloman, S. A., & Chaigneau, S. E. (2005). The HIPE theory of function. In L. C. E. v. d. Zee (Ed.), *Functional features in language and space:*

Insights from perception, categorization, and development (pp. 131–148). New York: Oxford University Press.

Berlin, B., Breedlove, D., & Raven, P. (1973). General principles of classification and nomenclature in folk biology. *American Anthropologist*, *74*, 214–242.

Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition*, *60*, 1–29.

Braisby, N. R., Franks, B., & Hampton, J. A. (1996). Essentialism, word use, and concepts. *Cognition*, *59*, 247–274.

Caramazza, A., & Shelton, J. R. (1998). Domain specific knowledge systems in the brain: The animate–inanimate distinction. *Journal of Cognitive Neuroscience*, *10*, 1–34.

Coley, J. D., Medin, D. L., & Atran, S. (1997). Does rank have its privilege? Inductive inferences within folkbiological taxonomies. *Cognition*, *64*, 73–112.

Diesendruck, G., Markson, L., & Bloom, P. (2003). Children's reliance on creator's intent in extending names for artifacts. *Psychological Science*, *14*, 164–168.

Dupré, J. (1981). Natural kinds and biological taxa. *The Philosophical Review*, *90*, 66–90.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Gelman, S. A. (2003). *The essential child*. New York: Oxford University Press.

Gelman, S. A., Coley, J. D., & Gottfried, G. (1994). Essentialist beliefs in children. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind* (pp. 341–365). Cambridge: Cambridge University Press.

Gelman, S. A., & Hirschfeld, L. A. (1999). How biological is essentialism? In D. L. Medin & S. Atran (Eds.), *Folkbiology* (pp. 403–446). Cambridge, MA: MIT Press.

Ghiselin, M. T. (1998). Folk metaphysics and the anthropology of science. *Behavioral and Brain Sciences*, *21*, 573–574.

Goodman, N. (1955). *Fact, fiction and forecast*. Cambridge, MA: Harvard University Press.

Gopnik, A., & Schulz, L. E. (Eds.). (in press). *Causal learning: Psychology, philosophy and computation*. Oxford: Oxford University Press.

Hadjichristidis, C., Sloman, S. A., Stevenson, R. J., & Over, D. E. (2004). Feature centrality and property induction. *Cognitive Science*, *28*, 45–74.

Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning Memory and Cognition*, *20*, 411–422.

Kalish, C. W. (1995). Essentialism and graded membership in animal and artifact categories. *Memory & Cognition*, *23*, 335–353.

Kalish, C. W. (1998). Natural and artificial kinds: Are children realists or relativists about categories? *Developmental Psychology*, *34*, 376–391.

Keil, F. C. (1989). *Concepts, kinds and cognitive development*. Cambridge, MA: MIT Press.

Keil, F. C. (1992). The origins of an autonomous biology. In M. R. Gunnar & M. Maratsos (Eds.), *Modularity and constraints in language and cognition, vol. 25. Minnesota symposium on child psychology* (pp. 103–138). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Keil, F. C. (1995). The growth of causal understanding of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary approach* (pp. 234–262). New York: Oxford University Press.

Keil, F. C. (2003). Folkscience: Coarse interpretations of a complex reality. *Trends in Cognitive Science*, *7*, 368–373.

Kelemen, D. (1999). Why are rocks pointy? Children's preference for teleological explanations of the natural world. *Developmental Psychology*, *35*, 1440–1452.

Kelemen, D., & Carey, S. (2007). The essence of artifacts: Developing the design stance. In S. Laurence & E. Margolis (Eds.), *Creations of the mind: Theories of artifacts and their representation*. Oxford: Oxford University Press.

Kripke, S. (1980). *Naming and necessity*. Oxford: Blackwell.

Kruschke, J. K. (1992). Alcove: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.

Lamberts, K., & Shapiro, L. (2002). Exemplar models and category-specific deficits. In E. Forde & G. W. Humphreys (Eds.), *Category-specificity in brain and mind* (pp. 291–314). Hove, UK: Psychology Press.

Lien, Y., & Cheng, P. W. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, *40*, 87–137.

Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition, 99*, 167–204.

Malt, B. C. (1990). Features and beliefs in the mental representation of categories. *Journal of Memory and Language*, *29*, 289–315.

Malt, B. C. (1994). Water is not H_2O. *Cognitive Psychology*, *27*, 41–70.

Malt, B. C. (1995). Category coherence in cross-cultural perspective. *Cognitive Psychology*, *29*, 85–148.

Malt, B. C., & Sloman, S. A. (2007). Artifact categorization: The good, the bad, and the ugly. In E. Margolis & S. Laurence (Eds.), *Creations of the mind: Essays on artifacts and their representation*. New York: Oxford University Press.

Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, *40*, 230–262.

Matan, A., & Carey, S. (2001). Developmental changes within the core of artifact concepts. *Cognition*, *78*, 1–26.

Mayr, E. (1982). *The growth of biological thought*. Cambridge, MA: Harvard University Press.

Medin, D. L., Coley, J. D., & Storms, G. H. B. (2003). A relevance theory of induction. *Psychonomic Bulletin and Review*, *10*, 517–532.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179–195). New York: Cambridge University Press.

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–393). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185–200.

Pauen, S. (2002). Evidence for knowledge-based category discrimination in infancy. *Child Development*, *73*, 1016–1033.

Pinker, S. (1997). *How the mind works*. New York: Norton.

Putnam, H. (1975). The meaning of 'meaning'. In K. Gunderson (Ed.), *Language, mind and knowledge* (pp. 131–193). Minneapolis: University of Minnesota Press.

Quine, W. V. O. (1970). Natural kinds. In N. Rescher (Ed.), *Essays in honor of Carl G. Hempel* (pp. 5–23). Dordrecht, The Netherlands: D. Reidel.

Quinn, P. C., & Eimas, P. D. (1996). Perceptual cues that permit categorical differentiation of animal species by infants. *Journal of Experimental Child Psychology*, *63*, 189–211.

Rakison, D. H., & Poulin-Dubois, D. (2001). Developmental origin of the animate–inanimate distinction. *Psychological Bulletin*, *127*, 209–228.

Rehder, B. (in press). Essentialism as a generative theory of classification. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy and computation*. Oxford: Oxford University Press.

Rehder, B., & Hastie, R. (2004). Category coherence and category-based property induction. *Cognition*, *91*, 113–153.

Rips, L. J. (2001). Necessity and natural categories. *Psychological Bulletin*, *127*, 827–852.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.

Scholl, B. J. (2005). Innateness and (Bayesian) visual perception: Reconciling nativism and development. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The structure of the innate mind* (pp. 34–52). Cambridge: Cambridge University Press.

Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology*, *40*, 162–176.

Shipley, E. F. (1993). Categories, hierarchies, and induction. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 30, pp. 265–301). San Diego, CA: Academic Press.

Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology*, *25*, 231–280.

Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology*, *35*, 1–33.

Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. New York: Oxford University Press.

Sloman, S. A., & Malt, B. C. (2003). Artifacts are not ascribed essences, nor are they treated as belonging to kinds. *Language and Cognitive Processes*, *18*, 563–582.

Sloman, S. A., & Over, D. E. (2003). Probability judgment from the inside and out. In D. Over (Ed.), *Evolution and the psychology of thinking: The debate* (pp. 145–169). Hove, UK: Psychology Press.

Sloman, S. A., & Rips, L. J. (1998). Similarity as an explanatory construct. *Cognition*, *65*, 87–101.

Sober, E. (1994). *From a biological point of view*. New York: Cambridge University Press.

Sterelny, K. (2003). *Thought in a hostile world: The evolution of human cognition*. Oxford: Blackwell.

Strevens, M. (2001). The essentialist aspect of naive theories. *Cognition*, *74*, 149–175.

Templeton, A. R. (1998). Human races: A genetic and evolutionary perspective. *American Anthropologist*, *100*, 632–650.