

Generics, prevalence, and default inferences

Sangeet Khemlani, Sarah-Jane Leslie, and Sam Glucksberg

{khemlani, sjleslie, samg}@princeton.edu

Departments of Psychology and Philosophy

Princeton University

Princeton, NJ 08540 USA

Abstract

We examine the extent to which people's judgments about whether a given member of a kind has a property are guided by prevalence estimates alone, or whether their acceptance of a generic generalization also impacts these judgments. Our data are not accounted for by prevalence estimates alone: acceptance of a generic disposes people to infer that, by default, a given member of a kind has the relevant property.

Keywords: Generics; default inferences; non-monotonic reasoning

Introduction

Generic sentences express generalizations, as in "ravens are black", "tigers are striped", and "mosquitoes carry the West Nile virus". These generalizations have a number of interesting features, for example, the generic "ducks lay eggs" is true, while "ducks are female" is false, yet the only ducks that ever lay eggs are the female ones. The generic "ticks carry Lyme disease" is true, even though the overwhelming majority of ticks do not carry the disease, yet "Canadians are right handed" is false, despite the fact that approximately 85% of Canadians are right handed.

Despite these complexities, generics are nonetheless a common part of everyday speech. Leslie (2007, 2008) has identified some features of these generalizations (see Table 1). The first concerns information that is characteristic of the kind in question. This notion is closely related to Prasada and Dillingham's (2006, 2009) notion of a principled connection, though it does not involve a commitment to the property in question being prevalent among members of the kind. Often, such characteristic information is prevalent amongst members of the kind, but this is not necessarily the case. Generics such as "ducks lay eggs", "pigs suckle their young", "lions have manes" and so on are examples of characteristic generics that are only true of a minority of the kind, i.e. of only the mature members of

one gender (Leslie, 2007, 2008). We will refer to these types of generics as 'minority characteristic' generics.

When a characteristic generic expresses a property that is prevalent among members of the kind, it can be thought of as expressing a principled connection in the sense of Prasada and Dillingham (2006, 2009). Examples of such principled characteristic generics include "tigers are striped", "ravens are black", "dogs have tails" and so on. It should be noted that such principled generics need not express exceptionless universal generalizations, since, e.g., *some* tigers are albino and therefore lack stripes. For the purpose of this paper, we will also identify a sub-kind of these principled characteristic generics, namely those that predicate properties that are true of *all* members of the kind without exception. Examples of these generics include "triangles have three sides" and "poodles are dogs". We will refer to these generics as 'quasi-definitional' generics.

Leslie's second main category of generics includes those that involve predication properties that are particularly striking, often potentially dangerous or damaging. When it comes to such properties – the sorts of properties that one would want to be forewarned about – very few members of the kind must have the property in question for the generic to be true. Consider, for example, "mosquitoes carry the West Nile virus", "sharks attack bathers", "ticks carry Lyme disease", "pit bulls maul children" and so on. These are all intuitively true generics, though it is only *very* few members of the kind that ever exhibit such properties. One might, however, wish to 'err on the side of caution' when it comes to such properties, however, and so be disposed to 'overgeneralize' them.

When it comes to information that is neither characteristic of the kind, nor particularly striking, then considerations of prevalence play more of a role in determining the truth of a generic. For example, generics such as "attics are dusty" and "cars have radios" would not be true if only a minority of

Table 1: Various types of generic generalizations

Predication type	Example	Truth Value of the Generic
Minority characteristic	Lions have manes	True
Principled	Dogs have tails	True
Quasi-definitional	Triangles have three sides	True
Striking	Pit bulls maul children	True
Majority	Cars have radios	True
High-prevalence True-as-Existentials	Canadians are right-handed	False
Low-prevalence True-as-Existentials	Rooms are round	False
False-as-existentials	Sharks have wings	False

attics was dusty or a minority of cars had radios. Leslie terms such statements ‘majority’ generics for obvious reasons. This category of generics/generalizations can be mapped onto Prasada and Dillingham’s (2006, 2009) category of *statistical connections*. We can thus identify a number of different types of generics: characteristic generics, including quasi-definitional, principled and minority characteristics, striking property generics, and majority generics.

There are examples of generics which are judged false, however, despite the property being had by a majority of the kind. For example, “books are paperbacks” and “Canadians are right-handed” seem false, yet most members of the kind have the properties in question. Leslie argues that such cases arise when the exceptions to the generalization have equally salient, concrete properties in place of the property being generalized. In order to accept “books are paperbacks” and “Canadians are right-handed” we have to ‘overlook’ the hardcover books and the left-handed Canadians (see Leslie, 2007, 2008 for detailed discussion).

In what follows, we refer to such statements as ‘high prevalence true-as-existentials’ or ‘high-prevalence TEs’, to reflect the fact that the statements are accepted in existential form (*some* Canadians are right-handed) but not in generic form. It is also easy to identify ‘low-prevalence TEs’ statements, e.g. “rooms are round”, “dogs are blind”, and “restaurants are Chinese restaurants”. We also make use of statements which are false even in existential form, such as “sharks have wings”. This latter class is referred as false-as-existential statements, or FEs (see Table 1 for a complete taxonomy, along with examples).

Leslie’s classification, as stated, is concerned almost exclusively with the circumstances under which generics are accepted or rejected. In this paper, we seek to address a related but different question: namely, how are people’s inferences affected by their belief in a generic generalization?

There is a body of work in AI and philosophical logic (e.g. Asher & Morreau, 1991, 1995; Pelletier & Asher, 1997; McCarthy, 1986; Reiter, 1987, and others) that has examined entailments involving generics, though much of it is formal/theoretical rather than empirical (for some notable exceptions to this trend, see Pelletier and Elio (1993) and Elio and Pelletier (1996)). The theorists in this tradition have also tended to focus on a restricted range of generics, namely the principled characteristic generics described above, with occasional inclusion of some majority generics. Striking generics have been almost entirely neglected in this literature.

Theorists in this tradition argue that minority characteristic generics such as “ducks lay eggs” and “lions have manes” are understood as implicitly restricted to the relevant gender. We understand, for example, “ducks lay eggs” to mean *female ducks lay eggs* (Pelletier & Asher, 1997; Asher & Morreau, 1995, and others). If this is how these sentences are understood, then they could be absorbed

into the category of principled characteristics, since it’s plausible that laying eggs is highly prevalent among female ducks. However, Leslie (2007, 2008) argues that this proposal faces philosophical and linguistic difficulties. An interesting upshot of the experiment we report here is that it provides empirical evidence against this view. The minority characteristic generics express generalizations concerning the entire kind, even though they are made true by only a minority of the kind.

In a series of collaborative articles, Asher, Morreau, and Pelletier incorporate insights from the AI literature into a philosophically and linguistically sensitive treatment of generics and their entailments (Asher & Morreau, 1991, 1995; Pelletier & Asher 1997). They discuss a number of entailment patterns involving generics, including ‘defeasible modus ponens’, exemplified by (1):

(1) Tigers have stripes

Tiggy is a tiger

Thus, Tiggy has stripes

They note that this argument, while not classically valid, is nonetheless compelling, and argue that this is because accepting a generic licenses one to infer that *by default* a given member of the kind has the generalized property. Their discussion, however, is more theoretical than empirical, and focuses almost exclusively on arguments that involve, in our terminology, principled characteristic generics, as is the case with (1). In our study, we investigate the extent to which people make such default inferences when the relevant generic cannot be classified as a principled characteristic generic. For example, will people be inclined to draw such inferences as those illustrated by (2) and (3)?

(2) Mosquitoes carry the West Nile virus

Buzzy is a mosquito

Thus, Buzzy carries the West Nile virus

(3) Ducks lay eggs

Quacky is a duck

So, Quacky lays eggs

We have previously investigated the conclusions people are willing to draw when presented with syllogisms involving generics (Khemlani et al., 2008, in preparation), but here we will study people’s inferences *outside* of a syllogistic context. In this way, our study constitutes a strong test of the inferential power of generics. Our participants were not presented with any generics whatsoever; rather, they were simply told that an individual was a member of a particular kind, for example,

Suppose you are told: Quacky is a duck

They were then asked to evaluate another non-generic statement by drawing on their own world knowledge:

What do you think of the following statement:

Quacky lays eggs

The experiment asked participants to judge whether and to what extent they were confident the above was true, confident it was false, or could not tell. We were particularly interested to learn whether these judgments

would be driven solely by people's beliefs about how prevalent the property is among members of the kind, or whether acceptance of the implicit background generic would have an impact on these inferences. If accepting a generic does, indeed, dispose one to infer that, by default, a given member of the kind has the property in question, then we would expect people to endorse these inferences more often when the implicit generic is true than when it is false. If this is so, then estimates of prevalence will fall short of explaining people's patterns of inference.

To test this, we intermixed items for which the implicit generic is false, yet the property in question is prevalent among members of the kind – i.e. high-prevalent true-as-existentials (high-prevalence TEs). For example, participants might be told “Joe is a Canadian” and asked to evaluate “Joe is right-handed”. These items are particularly interesting because the estimated prevalence of the property among the kind is high (i.e. most Canadians are right-handed), even though the generic is false. This experiment thus provides a direct test of the extent to which such inferences are driven by prevalence, as opposed to the background belief in the truth of the implicit generic. If the truth of the generic was found to override the estimated prevalence in guiding people's judgments, this would constitute very strong evidence for the power of the inferential profile of generics.

Method

Participants. 29 volunteers participated in the study through the Internet, and were recruited through Amazon's Mechanical Turk system for human interface tasks. None had any background in logic or computer science. They completed the experiment online using an interface written in Ajax. An independent sample of 17 online volunteers provided prevalence estimates (i.e., what percentage of Xs are Ys?) for each of the predication types that were used in the main experiment. These data appear in Table 2, below.

Materials. The items in the main phase of the experiment consisted of two statements: a statement about an individual's membership in a category (e.g., “Buzzy is a mosquito”) and a statement about the property that the member possessed (e.g., “Buzzy carries malaria”). Participants were asked to judge their level of confidence in the truth of that second statement. They registered their answers by selecting from a 7-point Likert scale (3 = I'm confident it's true, 0 = I couldn't possibly tell, -3 = I'm confident it's false). The statements concerning a property possessed by the individual expressed one of the eight different types of predication listed in Table 2. The last three types of predication – the high and low prevalence TEs and the FEs -- are all false in generic form, and so would allow us to determine the role that the truth of the implicit generic plays in these inferences. The names of the arbitrary members were selected so as to reveal no other relevant information in the context of the problem (e.g. for items in which gender was relevant, care was taken to use only gender-neutral names).

Design and procedure. The experiment was composed of two phases. In the main phase, participants completed the problems described above, in which they were asked to provide confidence ratings that an arbitrary member of a category possessed a particular property. The eight types of predication were distributed such that participants received 10 of each type except for 20 FE statements. The FE statements were included as a manipulation check to see whether participants were responding sensibly, and also served as a way to balance the frequency of responses across the ends of the rating scales. TE statements were classified as high-prevalence or low-prevalence by the prevalence estimates we had obtained in the norming study mentioned above. Participants also received two practice statements in the beginning of the study to familiarize themselves with the response scale, and the problems were presented in a different randomized order to each participant.

To help assess the impact of accepting a generic on these inferences, an auxiliary phase asked participants to judge whether a particular statement in generic form was true or false for the minority characteristic, striking, low-prevalence TE, high-prevalence TE, and majority predication. (We omitted the quasi-definitional, principled characteristics, and FEs from this phase, since acceptance of the first two and rejection of the third are at ceiling.) Accordingly, participants in the main phase who were asked to judge how confident they were that Buzzy carries malaria were asked in the auxiliary phrase to agree or disagree with the generic, “Mosquitoes carry malaria.” Half the participants received the main phase first and the auxiliary second, and half received the opposite order so as to determine whether participants who agreed with the generic were more likely to confer the property to an arbitrary member of the category. Participants received 10 generics of each of the five types of predication used in the auxiliary phase, and the generics were embedded within a set of 27 true and 27 false filler statements such as “eleven is a prime number” (true) and “tea is a soft drink” (false). They were told that the experiment tested general knowledge, and were instructed to press a designated button if they agreed or disagreed with each statement.

Results and discussion

Performance on the main phase of the experiment can be analyzed by coding participants' responses such that the points of the Likert scale are mapped to numerical values. Accordingly, when a participant was maximally confident of the truth (or falsity) of an arbitrary predication, their response was coded as +3 (or -3). In the context of the main task, the numerical values of the Likert scale loosely correspond to a measurement of inferential strength. That is, people make an inference when they attribute a property to an arbitrary member of a kind, and their level of confidence should predict their willingness to make this inference.

Table 2 shows the mean inferential strength attributed to various predication types, along with their corresponding

prevalence estimates. Order of presentation had no effect, and so these data were subjected to a one-way ANOVA, which revealed that inferential strength varied significantly by the different types of predication, $F(7,196) = 217.79$, $p < .0001$, $\eta_p^2 = .88$. Simple effects comparisons revealed that inferential strengths for principled characteristic predicates ($M = 2.6$) were not reliably larger than for quasi-definitional predicates ($M = 2.5$), $t(51.47)$, $p = .35$, $d = .25$; quasi-definitional predicates yielded reliably larger inferential strengths than minority characteristic predicates (1.7), $t(52.39)$, $p < .0001$, $d = 1.43$; inferential strengths for

Table 2: Mean inferential strengths and corresponding prevalence estimates for the different types of predication (proportion of true responses in parentheses)

Predication type	Inferential Strength	Prevalence Estimates
Principled	2.6 (97)	92
Quasi-definitional	2.5 (90)	*
Minority characteristic	1.7 (74)	64
Majority	1.5 (72)	*
Striking	0.7 (53)	33
High-prevalence TEs	0.6 (36)	*
Low-prevalence TEs	-0.3 (13)	17
False-as-existentials	-2.5 (4)	*
		5

Note: Asterisks between measures of inferential strength indicate a statistically significant difference in ratings between the two predication types.

minority characteristic predicates were not reliably larger than for majority predicates ($M = 1.5$), $t(54.49)$, $p = .45$, $d = .20$; inferential strengths for majority predicates were significantly larger than for striking predicates ($M = 0.7$), $t(51.11)$, $p < .005$, $d = .86$; inferential strengths for striking predicates were not reliably larger than high-prevalence TE predicates ($M = 0.6$), $t(44.77)$, $p = .63$, $d = .13$; inferential strengths for high-prevalence TE predicates were reliably larger than for low-prevalence TE predicates ($M = -.29$), $t(55.57)$, $p < .0001$, $d = 1.37$; and inferential strengths for low-prevalence TE predicates were reliably larger than for FE predicates ($M = -2.53$), $t(51.64)$, $p < .0001$, $d = 3.73$.

With respect to the influence of estimated prevalence, it is particularly informative that the inferential strength of striking predicates was not significantly different than high-prevalence TE predicates, while the estimated prevalences of the predicates diverged markedly, $t(18.72) = 6.09$, $p < .0001$, $d = 2.22$.

The table also presents (in parentheses) the proportion of true responses, i.e., whenever a participant's response was +3, +2, or +1. True responses indicate how often participants thought it was true that an arbitrary individual possessed the property in question. For instance, the third line of Table 2 shows that participants accepted statements like "Quacky lay eggs" 74% of the time, upon being told that Quacky is a duck.

These data suggest that principled properties and properties that hold of a kind by definition were attributed to

arbitrary members with a high degree of confidence. Properties that are characteristic of the kind but nevertheless are only possessed by one gender were also attributed to arbitrary members with a high degree of confidence, as were non-characteristic properties that hold for a majority of the kind. Participants were more wary about attributing the property to an arbitrary member of the kind when the property was striking or true only in existential form, though it is notable that the inferential strength of the striking predication was the same as for the high prevalence TEs, despite the much lower prevalence estimates for the striking properties. The low prevalence TEs were more comparable to the striking in terms of prevalence estimates, but were accorded significantly lower inferential strength. Finally, participants were confident that an arbitrary member of a kind should not be attributed a property that holds for no member of the kind.

Analysis of generic acceptance. Participants' endorsement of a generic statement may drive their tendency to attribute a given property to an arbitrary member of a class. That is, people who agree with the statement "ticks carry Lyme disease" and subsequently learn that Jumpy is a tick may be more likely to infer that Jumpy carries Lyme disease than those who do not agree with the generic. The auxiliary phase of the experiment offered insight into this issue by presenting participants with generics corresponding to the kinds and predicates they experienced in the main phase of the experiment.

Participants agreed with minority characteristic generics on 85% of the trials, majority generics on 80% of the trials, and striking generics on 75% of the trials; these results replicate earlier findings (Khemlani et al., 2007). Participants agreed with high-prevalence and low-prevalence TEs on 43% and 20% of the trials respectively. There was no effect of presentation order, and so these data were subjected to a one-way ANOVA, which yielded a main effect of predication type, $F(4, 112) = 56.13$, $p < .0001$, $\eta_p^2 = .67$.

To analyze whether participants' endorsement of a generic statement affected their tendency to attribute the property to an arbitrary member of a kind, mean inferential strengths can be separated by whether participants accepted or rejected the generic statement. These data are summarized in Table 3, and again the proportions of true responses (i.e., +1, +2, or +3) are included in parentheses. Pairwise comparisons between whether they agreed or disagreed with a generic were conducted for each predication type in the main phase. Note that very few

Table 3: Mean inferential strengths separated by predication types and by acceptance of the generic (proportion of true responses in parentheses)

Predication type	Accepted generic	Rejected generic
Minority characteristic	2.0 (85)	-0.4 (23)
Majority	1.7 (78)	0.7 (47)
Striking	1.1 (65)	-0.6 (15)
High-prevalence TEs	1.0 (52)	0.3 (23)
Low-prevalence TEs	0.2 (21)	-0.4 (12)

participants disagreed with minority characteristic or majority generics, and likewise participants tended not to accept low-prevalence TE generics. Statistical comparisons between agreement vs. disagreement of a generic for minority characteristic, majority, and low-prevalence TEs should thus be viewed critically (these proportions are italicized in Table 3). Nevertheless, pairwise comparisons yielded significant effects of generic agreement for all types of predication, including minority characteristic, $t(20.21) = 7.75$, $p < .0001$, $d = 2.51$; majority, $t(23.25) = 2.33$, $p < .05$, $d = .77$; striking, $t(32.80) = 5.43$, $p < .0001$, $d = 1.64$; high-prevalence TE, $t(37.83) = 4.81$, $p < .0001$, $d = 1.30$; and low-prevalence TE predication, $t(27.48) = 2.37$, $p < .05$, $d = .73$.

These findings reinforce the idea that accepting the generic has a substantial effect on these inferences. When we separate the participants who accepted the generic from those who did not, we find that the tendency to endorse the inference in question becomes even greater. For example, those participants who accepted a given striking generic endorsed the attribution of the relevant property to a given member of the kind 65% of the time – i.e., participants who accepted “ticks carry Lyme disease” agreed with the claim that Jumpy carries Lyme disease a stunning 65% of the time.

General discussion

This experiment offers a dramatic illustration of the inferential power of generics. It is interesting to note that, among the different types of predication that are true as generics, prevalence estimates were a good predictor of the extent to which participants endorsed the inference. Similarly, among the types of predication that are *false* as generics, prevalence estimates were again a good predictor. However, if we consider all the predication types together – without separating them out based on the truth/falsity of the generic form – then prevalence falls short of explaining the data. The striking predication produced a comparable amount of inferences to the high-prevalence TEs, yet the prevalence estimates for the striking predication were far lower than those for the high-prevalence TEs. The high prevalence TEs produced significantly fewer inferences than the minority characteristics, yet the prevalence estimates for the two were not reliably different. Thus it seems that the truth of the implicit background generic produced an inferential ‘boost’, above and beyond the estimates of prevalence. This suggests that the acceptance of a generic does indeed dispose people to infer by default that a given instance has the relevant property.

It is telling to compare the inferential power of the striking predication to those of the high-prevalence TEs. Despite the very different prevalence estimates – 33% versus 60% – participants accorded these two types of predication the same degree of inferential power. The prevalence estimates for the high-prevalence TEs were more comparable to those for the minority characteristic and

majority predication, but these latter types of items were accorded far more inferential power.

Finally, we wish to draw attention to an interesting and telling aspect of our data. As we discussed in the introduction, many theorists assume that minority characteristic generics can be analyzed as generalizing over only one gender, or some other comparable restriction. For example, “Ducks lay eggs”, they argue, is really understood to mean *female ducks lay eggs* (Pelletier and Asher, 1997; Asher and Morreau 1991, 1995 and others). Our data here constitute an empirical refutation of this popular line of thought, even by the lights of the proponents of this view.

The idea here is simple. If people understand “ducks lay eggs” to mean *female ducks lay eggs* – if that is how the generic is understood and represented – then there should be no temptation to conclude that an arbitrary duck lays eggs. One would only judge that a given duck lays eggs to the same (limited) extent that they would judge that a given duck is female. The participants were decidedly not tempted to judge that arbitrary animals belonged to a particular gender to the extent they were tempted to judge that a given duck lays eggs – we had included some such items, including “ducks are female”, among our high prevalence TEs in the present experiment. The responses to these ‘gender-ascribing’ TEs were, if anything, lower than the overall responses to the high prevalence TEs in general. Yet we found that people did not hesitate to conclude that a given duck lays eggs, and similarly for the other minority characteristics. This inferential disposition is explicable only if one posits that people understand “ducks lay eggs” to be an unrestricted generalization across the entire kind *duck*. Indeed, Pelletier and Asher – who espoused a gender-restricted view – explicitly note that an experiment such as ours would constitute a decisive test of this hypothesis. They write:

One way to test this hypothesis [of gender-restriction] is to look at how speakers treat such problematic generics as [*ducks lay eggs*] when they exploit them in inferences. We have argued that an inference like Defeasible Modus Ponens is defeasibly valid and so should be part of a theory of generics. But many speakers express reluctance to draw even the defeasible conclusion that if Allie is a duck and ducks lay eggs, then Allie lays eggs. It would depend, they say, upon whether Allie is a male or a female duck (Pelletier & Asher, 1997, pp. 1166-1167).

We failed to find any such reluctance. Participants overwhelmingly endorsed the relevant inferences. Based on the minimal prompt that Quacky is a duck, participants who accepted the generic “ducks lay eggs” agreed 83% of the time that Quacky lays eggs. If generics have powerful inferential profiles, and the generic “ducks lay eggs” concerns duck-kind in general – and not just the female ones – then we would predict exactly this result.

The study described here demonstrates that generics have a powerful inferential profile which, to some extent, trumps

estimates of prevalence in guiding people's default inferences. This is all the more remarkable because people's average prevalence estimates were already inflated relative to the facts, at least for the striking and minority characteristic predication.

One direction in which we plan to develop these findings is to investigate the extent to which they may apply to social categorization and stereotyping. Leslie (in press) argues that some prejudiced generalizations, such as *Muslims are terrorists*, may be cognitively similar to striking property generalizations such as the socially innocuous *ticks carry Lyme disease*. If this is so, then our results shed light on the troubling inferential power of such judgments.

Acknowledgments

This study is based on work supported in part by a National Science Foundation Graduate Research Fellowship awarded to the first author. We thank Jeremy Boyd, Adele Goldberg, Phil Johnson-Laird, Mark Johnston, Olivia Kang, and Lance Rips for their many helpful discussions and suggestions.

References

Asher, N. & Morreau, M. (1991). Common sense entailment: A modal theory of non-monotonic reasoning. *Proceedings of the 12th IJCAI*, 387-392. San Mateo, CA: Morgan Kaufman.

Asher, N. & Morreau, M. (1995). What some generic sentences mean. In G. Carlson and F. J. Pelletier (eds.) *The generic book*. Chicago: Chicago University Press.

Elio, R. & Pelletier, F.J. (1996). On reasoning with default rules and exceptions. In G. Cottrell (Ed.): *Proceedings of the 18th Annual Conference of the Cognitive Science*. Hillsdale, NJ: Lawrence Erlbaum.

Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. New York: Oxford University Press.

Khemlani, S., Leslie, S.J., Glucksberg, S., & Rubio-Fernandez, P. (2007). Do ducks lay eggs? How people interpret generic assertions. In: *Proceedings of the 29th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Khemlani, S., Leslie, S.J., & Glucksberg, S. (2008). Syllogistic reasoning with generic premises: The generic overgeneralization effect. In: *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Khemlani, S., Leslie, S.J., & Glucksberg, S. (in preparation). Reasoning with generics.

Leslie, S.J. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, 21 (1).

Leslie, S.J. (2008). Generics: Cognition and acquisition. *Philosophical Review*, 117 (1).

Leslie, S.J. (in press). The original sin of cognition: Fear, prejudice, and generalization. *The Journal of Philosophy*.

Medin, D.L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (eds.) *Similarity and Analogical Reasoning*. New York: Cambridge University Press.

McCarthy, J. (1986). Applications of circumscription to formalizing commonsense knowledge. *Artificial Intelligence*, 28.

Pelletier, F. & Asher, N. (1997). Generics and defaults. In J. van Benthem and A. ter Meulen (eds.) *Handbook of Logic and Language*. Cambridge, MA: MIT Press.

Pelletier, F. & Elio, R. (1993) Human benchmarks on AI's benchmark problems. In: *Proceedings of the 15th Congress of the Cognitive Science Society*. Boulder, Co: Cognitive Science Society.

Prasada, S. & Dillingham, E. (2006). Principled and statistical connections in common sense conception. *Cognition*, 99 (1).

Prasada, S. & Dillingham, E. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, 33.

Reiter, R. (1987). Nonmonotonic reasoning. In *Annual Review of Computer Science*, ed. by J. F. Traub, N. J. Nilsson, and B. J. Grosz. Palo Alto: Annual Reviews Inc.