Argument Scope in Inductive Reasoning:
Evidence for an Abductive Account of Induction

Samuel G. B. Johnson[1], Thomas Merchant[2], & Frank C. Keil[1]
(samuel.johnson@yale.edu, thomas_merchant@brown.edu, frank.keil@yale.edu)
[1]Dept. of Psychology, Yale University, 2 Hillhouse Ave., New Haven, CT 06520 USA
[2]Dept. of Cognitive, Linguistic, and Psychological Sciences, Brown University, 190 Thayer St., Providence, RI 02912 USA

Abstract

Our ability to induce the general from the specific is a hallmark of human cognition. Inductive reasoning tasks ask participants to determine how strongly a set of premises (e.g., Collies have sesamoid bones) imply a conclusion (Dogs have sesamoid bones). Here, we present evidence for an *abductive theory* of inductive reasoning, according to which inductive strength is determined by treating the conclusion as an explanation of the premises, and evaluating the quality of that explanation. Two inductive reasoning studies found two signatures of explanatory reasoning, previously observed in other studies: (1) an *evidential asymmetry* between positive and negative evidence, with observations casting doubt on a hypothesis given more weight than observations in support; and (2) a *latent scope effect*, with ignorance about potential evidence counting against a hypothesis. These results suggest that inductive reasoning relies on the same hypothesis evaluation mechanisms as explanatory reasoning.

Keywords: Inductive reasoning; abductive inference; explanation; scope; hypothesis evaluation.

Introduction

Learning often requires us to induce the general from the specific. Yet, David Hume (1977/1748) famously noted that inductive inferences are never logically valid, because subsequent observations could falsify them. Thus, cognitive scientists have long recognized that humans have both deductive and inductive capacities—we can reason deductively, where known information implies with certainty the truth of the conclusion, and we can reason inductively, where known information implies only with some probability the truth of the conclusion.

To study inductive reasoning, psychologists typically confront participants with arguments like:

> *Collies have sesamoid bones.*
> *Therefore, dogs have sesamoid bones.*

People are asked to rate the extent to which the premise(s) (here, Collies having sesamoid bones) support the conclusion (Dogs having sesamoid bones). This argument can be compared to the following (better) argument:

> *Collies have sesamoid bones.*
> *Poodles have sesamoid bones.*
> *Therefore, dogs have sesamoid bones.*

This argument is superior because it has an additional supporting premise—a phenomenon known as *premise monotonicity* (Osherson et al., 1990).

Several theoretical models have arisen from the results of such tasks (Heit, 2000). Some models (e.g., Osherson et al., 1990; Sloman, 1993) explain inductive reasoning in terms of similarity. For example, the Osherson et al. (1990) model explains premise monotonicity by observing that the total similarity of the premise categories to the conclusion category is greater when there are more premises, because more premises can cover a larger part of the similarity space. Hume (1977/1748) took a similar approach. His *Principle of the Uniformity of Nature* proposes that the unknown will resemble the known: That is, the more similar the unknown conclusion is to the known premises, the stronger the inference from premise to conclusion.

Other researchers have posited more active, flexible cognitive processes, arguing that similarity alone is too unconstrained to account for all of induction. For example, people rely on different aspects of premise–conclusion similarity, depending on the property (Heit & Rubinstein, 1994). When the property is anatomical, more anatomically similar premise categories confer greater strength to their conclusions (e.g., if the conclusion is that a bat has a certain type of liver, then a mouse sharing that property is better evidence for the conclusion compared to a sparrow). Conversely, when the property is behavioral, more behaviorally similar premise categories confer greater strength to their conclusions (e.g., if the conclusion is that a bat never travels in the same direction as the sun, then a sparrow sharing that property is better evidence for the conclusion compared to a mouse). Thus, similarity alone does not determine inductive strength, but must be combined with other critical assumptions.

Recently, Bayesian theories of inductive reasoning have emerged as a possible way to flesh out these assumptions (Heit, 1998; Kemp & Tenenbaum, 2009). According to these theories, people approach inductive reasoning problems by estimating the probability of the conclusion, given the truth of the premises. Bayesian theories are situated at the computational level and posit normative probability calculations. This assumption generally works in favor of such theories, because most inductive reasoning phenomena appear to be normative—for instance, more positive evidence in the form of additional premises usually *does* make the conclusion more likely to be true. But this normativity would be a liability if people sometimes make non-normative inductive inferences.

Here, we argue for a third, *abductive* position. Peirce (1997/1903) distinguished between enumerative or Humean induction—becoming increasingly confident in a generalization as instances of it accumulate—and *abduction*—becoming confident in a generalization to the

extent that it is a good *explanation* of the data (see also Lipton, 2004 and McDonald, Samuels, & Rispoli, 1996). More precisely, in an abductive inference, the reasoner begins with a set of observations, thinks of potential explanations that would account for the observations, and accepts an explanation to the extent that it is satisfactory according to a set of "explanatory virtues." In a similar fashion, we propose that reasoners approach inductive reasoning tasks by treating the premises as evidence, and the conclusion as a potential explanation to be evaluated.

This account is in the general spirit of Bayesian accounts, with the caveat that computational-level Bayesian accounts claim only that participants calculate posterior probabilities, without any further specificity as to what reasoning processes lead to those posteriors at the algorithmic level. Recent evidence suggests that people use a set of explanatory heuristics or virtues to evaluate explanations (Johnson, Jin, & Keil, 2014; Johnson, Rajeev-Kumar, & Keil, 2014, 2015; Johnston, Johnson, Koven, & Keil, 2015; Lombrozo, 2007). These heuristics, while computationally tractable and likely to be approximately normative under favorable conditions, can lead to systematic errors under less favorable (laboratory) conditions. In this paper, we look for two signatures of abductive inference to support the position that inductive reasoning is a species of explanatory reasoning.

Consider again the conclusion that dogs have sesamoid bones. This conclusion (or explanation, in our abductive terms) would make a variety of predictions. For example, it would predict that German Shepherds have sesamoid bones. The set of predictions made by an explanation is known as its *scope*, and an explanation's scope can be divided into three types of evidence—positive, negative, and latent (Johnson, Johnston, Toig, & Keil, 2014). From a normative perspective, positive evidence (e.g., German shepherds having sesamoid bones) should count in favor of an explanation by confirming a prediction; negative evidence (German shepherds *not* having sesamoid bones) should count against an explanation by disconfirming a prediction; and latent evidence (e.g., not knowing whether or not German shepherds have sesamoid bones) counts neither for nor against an explanation because it is simply unavailable. Prior research has identified two phenomena in explanatory reasoning concerning these evidence types, which we treat as signatures of abductive inference—an asymmetry between positive and negative evidence (Johnson & Keil, 2015), and an aversion to explanations positing latent evidence (Khemlani, Sussman, & Oppenheimer, 2011).

First, people do not treat confirmed predictions (positive evidence) and disconfirmed predictions (negative evidence) symmetrically. In studies of causal explanation (Johnson & Keil, 2015), an explanation that posited two observed effects was seen as somewhat better than an explanation that posited only one observed effect. However, an explanation that posited one observed effect was seen as far better than an explanation that posited one

observed and one disconfirmed effect. That is, negative evidence is weighed more than positive evidence—a phenomenon reminiscent of loss aversion in decision-making (Kahneman & Tversky, 1979).

Second, people do not treat unverified predictions as irrelevant, but usually count them as evidence *against* an explanation. When comparing an explanation that makes one confirmed prediction (*Y*) against an explanation that makes one confirmed prediction as well as a latent prediction of unknown truth (*YD*), people prefer the explanation that does not make the latent prediction, even if the explanations have equal prior probability (Khemlani et al., 2011; see also Johnson, Rajeev-Kumar, & Keil, 2014 for the mechanism). This bias is non-normative because the explanations are equally likely if they have equal prior probabilities. If found in inductive reasoning, this phenomenon would be incompatible with Bayesian accounts that posit normative probabilistic reasoning.

We test for these explanatory signatures in two experiments. In both experiments, we measure the effects of adding positive, negative, and latent evidence. We predicted that negative evidence would count against a conclusion more than positive evidence would count in its favor. We also predicted that people would show a latent scope bias, rating conclusions as weaker when some of their potential evidence is latent. Experiment 1 measured probability judgments (facilitating comparison with Bayesian accounts) and Experiment 2 measured argument strength (facilitating comparisons with other empirical studies of inductive reasoning).

Experiment 1

In Experiment 1, participants evaluated four different types of arguments. Participants read a cover story, stating that the premises were generated by an expert on the topic and that the judgments were based on experimental evidence. This was done so as to minimize concerns about possible pragmatic inferences (e.g., that ignorance about a premise signaled that the premise was false) by giving non-pragmatic justifications for why the expert had the information that they had.

The baseline for comparison were arguments consisting of two premises and a conclusion. We referred to this argument type as *YY*, because both premises were positive. For example, the premises for one YY item read:

A study found that rainbow trout have T-A enzymes.
A study found that brown trout have T-A enzymes.

The conclusion for this (and following) examples was:

Fish have T-A enzymes.

To compare the relative effects of positive and negative evidence, matched versions were created that included either an additional positive premise or an additional negative premise (but the same conclusion). A *YYY* item, with additional positive evidence, read as follows:

A study found that rainbow trout have T-A enzymes.
A study found that brown trout have T-A enzymes.
A study found that clownfish have T-A enzymes.

The corresponding *YYN* item, with additional negative evidence, read as follows:

> *A study found that rainbow trout have T-A enzymes.*
> *A study found that brown trout have T-A enzymes.*
> *A study found that clownfish do <u>not</u> have T-A enzymes.*

Note that the third premise in both cases adds to the *premise diversity*—that is, results in a more varied set of premises to project to the conclusion. This ought to lead to relatively large effects of positive evidence, since premise diversity is among the most robust findings in category-based induction (Osherson et al., 1990), working against our hypothesis of finding a larger effect of negative than of positive evidence.

Finally, to test for an effect of latent scope, a fourth argument (*YYD*) was included, wherein the truth of the additional premise was unknown and hence in the latent scope of the conclusion category.

> *A study found that rainbow trout have T-A enzymes.*
> *A study found that brown trout have T-A enzymes.*
> *We do <u>not</u> know if clownfish have T-A enzymes, because the study results have not yet come back from the lab.*

If participants are evaluating the argument by evaluating whether the conclusion is a good explanation of the premises, then this YYD argument should be less convincing than the YY argument, even though the additional premise in fact makes the conclusion no more or less probable, normatively speaking.

Method

We recruited 200 participants from Amazon Mechanical Turk for Experiment 1; 6 were excluded from analysis because they incorrectly answered more than 30% of a set of multiple choice check questions.

Each participant completed four items—one each in the YY, YYY, YYN, and YYD conditions, formatted as above. For each item, participants read the premise and conclusion statements, and responded to the question "Assuming that the premises are all true, how likely do you think it is that the conclusion is also true?" on a scale from 0 ("Very unlikely") to 10 ("Very likely"). These items were drawn from four domains—social kinds, artifacts, living natural kinds, and non-living natural kinds. Condition and domain were balanced using a Latin square. Two different sets of items were used (e.g., one biological item concerned fish, as above, and another concerned plants), and participants were randomly assigned to one set of items. Items were completed in a random order, but the premises were always listed in the same order (as in the above examples).

Results and Discussion

As shown in Figure 1, participants' opinions about the conclusion probabilities depended on the nature of the premises. Compared to the YY arguments, the YYY arguments conferred more certainty on their conclusions,

indicating an effect of adding positive evidence [$M = 6.84$, $SD = 2.63$ vs. $M = 5.81$, $SD = 2.80$; $t(193) = 4.24$, $p < .001$, $d = 0.30$]. Similarly, compared to the YY arguments, the YYN arguments conferred *less* certainty on their conclusions, indicating an effect of adding negative evidence [$M = 2.91$, $SD = 2.75$ vs. $M = 5.81$, $SD = 2.80$; $t(193) = -12.48$, $p < .001$, $d = -0.90$]. These results accord with previous studies demonstrating effects of positive (Osherson et al., 1990) evidence in inductive reasoning, and go beyond those studies in showing a robust effect of negative evidence.

These effects are qualitatively consistent with many models of inductive reasoning, including similarity, Bayesian, and abductive approaches. However, the relative size of the effects is key. The effects of negative and positive evidence were not symmetric: Negative evidence had a far more dramatic effect on the perceived probability compared to positive evidence [$M = 2.90$, $SD = 3.24$ vs. $M = 1.03$, $SD = 3.37$; $t(193) = 4.55$, $p < .001$, $d = 0.33$]. This result is consistent with work on abductive inference, showing that in general negative evidence has a greater deleterious effect on explanatory judgments compared to the advantageous effect of positive evidence (Johnson & Keil, 2015). Although other accounts of inductive reasoning (e.g., Bayesian approaches) may have resources to explain this asymmetry, this result seems to favor an abductive account. Further, this asymmetry cannot be the result of asymmetric floor and ceiling effects, because the YYN ratings were actually closer to the floor than the YYY ratings to the ceiling.

The most compelling evidence in favor of the abductive account, however, concerns the effect of adding a premise espousing ignorance. The YY arguments were seen as more convincing than the YYD arguments, indicating an effect of latent scope [$M = 5.29$, $SD = 2.51$ vs. $M = 5.81$, $SD = 2.80$; $t(193) = 2.32$, $p = .021$, $d = 0.17$]. Rather than ignoring the unknown premise, participants counted it as evidence against the conclusion. This result is easily predicted by the abductive account of inductive reasoning, since people consistently find explanations with wide latent scope (i.e., making unverified predictions) less satisfying and less probable than explanations with narrow latent scope (Johnson, Rajeev-Kumar, & Keil,
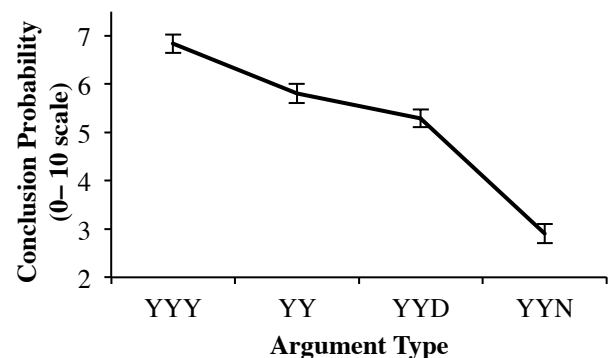

Figure 1: Results of Experiment 1. Bars represent 1 SE.

2014; Khemlani et al., 2011). This result is more difficult to explain with similarity or Bayesian accounts.

Experiment 2

Experiment 2 built off of Experiment 1 in two ways. First, we included both essential properties (e.g., a kind of lawyer having a certain personality type) and accidental properties (a kind of lawyer using a certain type of office software). Because properties can vary in their projectibility to new categories (Heit, 2000), we sought to test whether evidential asymmetries and latent scope bias would also extend to less essential properties. In particular, if people see such properties as less diagnostic of category membership, they may be less likely to use explanatory reasoning to account for those properties. Second, rather than asking directly about the probability of the conclusion, as in Experiment 1, we asked about argument strength ("Please rate how well these premises support this conclusion"). This was done to make the results more comparable with previous studies of inductive reasoning, which often measure argument strength rather than probability (Osherson et al., 1990).

Method

We recruited 200 participants from Amazon Mechanical Turk for Experiment 2; 12 were excluded from analysis because they incorrectly answered more than 30% of a set of multiple choice check questions.

The procedure was identical to Experiment 1, except that participants completed both an essential item and an accidental item from each of the four domains (e.g., in the biological domain, the essential version of the fish item and the accidental version of the plant item, or the converse), for a total of 8 items. Argument strength was measured on 0 (premises support conclusion "very poorly") to 10 ("very well").

Results and Discussion

As shown in Figure 2, participants once again found the arguments to be of differing strength, depending on the nature of the premises. Somewhat unexpectedly, the effects did not differ as a function of whether the property was accidental or essential ($ts < 1$, $ps > .60$ for the interactions), so we collapse across this factor.

The results were similar to those of Experiment 1. First, the YYY arguments were rated stronger than the YY arguments [$M = 6.49$, $SD = 2.53$ vs. $M = 5.20$, $SD = 2.70$; $t(187) = 8.33$, $p < .001$, $d = 0.61$], and the YY arguments were rated stronger than the YYN arguments [$M = 3.38$, $SD = 2.57$; $t(187) = 8.14$, $p < .001$, $d = 0.59$], indicating effects of positive and negative evidence on argument strength, respectively. Once again, the effect of adding negative evidence was larger than the effect of adding positive evidence, although this trend did not reach significance [$M = 1.82$, $SD = 3.07$ vs. $M = 1.29$, $SD = 2.13$; $t(187) = 1.63$, $p = .104$, $d = 0.12$]. Second, the YY arguments were rated stronger than the YYD arguments
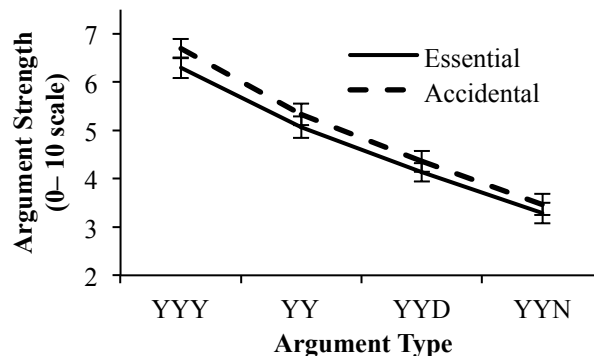

Figure 2: Results of Experiment 2. Bars represent 1 SE.

[$M = 5.20$, $SD = 2.70$ vs. $M = 4.25$, $SD = 2.49$; $t(187) = 4.90$, $p < .001$, $d = 0.36$], indicating a latent scope effect.

These results show that Experiment 1's results concerning the conclusion probability also extend to the more traditional measure of argument strength. Surprisingly, the effects of positive, negative, and latent evidence did not depend on the nature of the property—essential or accidental—despite previous demonstrations that properties vary in their projectibility (Heit, 2000). One possibility is that participants essentialized even the accidental properties, treating them as a critical part of the category. This possibility seems especially likely given the nearly identical means for both types of properties.

Alternatively, some participants could have interpreted the argument strength measure as asking what formal properties make for a good argument. They might then rely less on their intuitive judgments of probability, and more on their folk theories of argumentation (e.g., Corner & Hahn, 2009). This interpretation also may be consistent with the weaker asymmetry in Experiment 2 between positive and negative evidence (since people's folk theories of argumentation seem equally likely to weigh positive and negative evidence more heavily) and with the larger latent scope effect (since ignorance is often taken as a sign of poor argumentation; e.g., Durik, Britt, Reynolds, & Storey, 2008).

Of course, neither interpretation undermines our core claim that people use explanatory principles in inductive reasoning, since we also found these results with a less ambiguous dependent measure in Experiment 1. Nonetheless, future research might use other dependent measures (such as probability, plausibility, or explanatory judgments) to study the effects of positive, negative, and latent evidence given properties of varying projectibility.

General Discussion

Much of cognition consists of going beyond the known, to infer new knowledge—reasoning inductively. Here, we demonstrated two phenomena of inductive reasoning that speak in favor of an *abductive* account. According to this account, when people reason inductively from premises to a conclusion, they judge the conclusion to be supported by the premises to the extent that the conclusion appears

to be a good explanation of the premises (or a consequence of such an explanation). Based on recent research, we tested two signatures of abductive reasoning.

Evidential asymmetry. First, negative evidence in the form of disconfirmed predictions is usually weighed more heavily than positive evidence (Johnson & Keil, 2015), just as losses are weighed more heavily than gains in decision-making (Kahneman & Tversky, 1979). This same asymmetry was observed in inductive reasoning: In Experiment 1, negative premises counted more strongly against a conclusion than positive premises counted in its favor. Although this effect only reached marginal significance in Experiment 2, the consistent pattern across studies (and large mean difference in Experiment 1) leaves little empirical doubt about this result.

Might this result be accounted for in terms of similarity or probability? Similarity-based accounts assume that the premises support the conclusion to the extent that the premise categories 'cover' as much of the conclusion category as possible. To our knowledge, similarity-based accounts have not made specific predictions about negative evidence, but they could be extended to account for negative evidence leading to weaker argument strength: Negative evidence explicitly limits the coverage of the premise categories. However, if negative evidence is seen merely as the negation of potential positive evidence, then a similarity account would seem to predict symmetry between positive and negative evidence, contrary to our results.

An alternative possibility is that negative evidence should *completely* negate the conclusion, if the conclusions are interpreted as predicating a property to *all* members of a category (i.e., if clownfish do not have T-A enzymes, then clearly it is not the case that *all* fish have T-A enzymes). This possibility, however, is untenable in light of research on generic language (Leslie, 2008). Statements such as "Fish have T-A enzymes" allow for exceptions, just as explanatory generalizations allow for disconfirmatory evidence to be explained away by other factors (e.g., an experiment might fail not because a theory was wrong, but because the methodology was flawed). Further, people do not treat subordinate categories (e.g., clownfish) as inheriting all properties from their superordinate categories (fish)—this is clear both in our data (ratings for the *YYN* arguments were far above the floor) and in data from Sloman (1998).

Bayesian theories may be better able to account for this asymmetry. Although not predicted *a priori* by these accounts, people might think that there are more "innocent" ways (i.e., alternative explanations other than the hypothesis being false) for an explanation's prediction to be accidentally confirmed than for it to be disconfirmed—perhaps because an accidental confirmation requires only an unexplained generative cause, but an accidental disconfirmation requires an unexplained preventive cause, *in addition to* the explained cause. If so, negative evidence would be more diagnostic than positive evidence. This account is consistent with existing data—and with both the Bayesian and abductive accounts—but has not yet been explicitly tested.

Latent scope bias. Second, *latent* evidence in the form of unverified predictions is usually treated like negative evidence, and counted against a hypothesis (Khemlani et al., 2011). This effect is non-normative from a probabilistic standpoint, because observations of which we are completely ignorant are equally consistent with any hypothesis. Both Experiments 1 and 2 found this bias in inductive reasoning: Premises espousing ignorance about some possible evidence (e.g., "We do not know if clownfish have T-A enzymes, because the study results have not yet come back from the lab") were counted against the conclusion.

This result not only is predicted *a priori* by an abductive account of inductive reasoning, but is difficult to reconcile with similarity- or probability-based accounts. It is unclear why a premise asserting ignorance about a category should enter into the similarity computation at all, much less why it should count negatively. (Indeed, *positive* latent scope effects can be obtained when the base rate of an observation is very high, creating a further obstacle for a similarity-based account; Johnson, Rajeev-Kumar, & Keil, 2014). And the latent scope bias is even more difficult to square with probabilistic accounts, because it is non-normative.

Other accounts. If inductive inference is not driven by similarity or probability, then what *does* do the driving?

According to *relevance theory* (Medin, Coley, Storms, & Hayes, 2003), people perform inductive reasoning by assuming that the premises are informative with respect to the conclusion. Although it is not clear how this account would make sense of the explanatory asymmetry, it potentially provides an alternative explanation for the latent scope effect—that is, people might assume that an informative speaker would not flag their ignorance unless it was pragmatically relevant to interpreting the conclusions. However, in Experiment 1, an alternative (non-conversational) reason was given for ignorance (i.e., the results not yet being back from the lab), and a latent scope effect was still found. Other research has also found latent scope effects when pragmatic inferences are blocked (Johnson, Rajeev-Kumar, & Keil, 2014); thus, relevance theory cannot fully explain these results.

Of existing theories, these results are most consistent with *hypothesis-testing theory* (McDonald et al., 1996), which also holds that people treat the premises as evidence and the conclusion as a hypothesis or explanation. The current abductive model might be best seen as a way of fleshing out the theory of McDonald et al. (1996), in light of recent findings in abductive reasoning. Indeed, McDonald et al. report several findings consistent with our abductive model, most notably that arguments with more plausible alternatives to the conclusion are rated weaker than arguments with fewer such alternatives. As more alternative explanations are

made available, the target explanation (conclusion) becomes increasingly unlikely to be the most satisfactory.

Future theoretical and empirical work can be done to pinpoint abductive interpretations of other inductive reasoning phenomena (see Heit, 2000 for a review). For example, more diverse premises usually lead to more confidence in the conclusion (Osherson et al., 1990). One potential reason for this phenomenon is that alternative explanations of more diverse premises would need to be highly complex, and people are averse to complex explanations (Lombrozo, 2007). Boundary conditions on the simplicity preference (Johnson, Jin, & Keil, 2014) may be useful for empirically distinguishing this abductive account from competing theories.

Perhaps the greatest promise of an abductive theory is its potential to unify research on inductive reasoning with the growing body of research on explanatory reasoning throughout psychology. Psychological processes from categorization to vision to language understanding have been cast in terms of explanatory inference. Inductive inference may be a special case of a far more general process of explanatory reasoning that pervades much of cognition—a process guided by a set of fallible yet ordinarily truth-tracking heuristics, which can allow us to flexibly produce new knowledge from old.

Acknowledgments

References

Corner, A., & Hahn, U. (2009). Evaluating science arguments: Evidence, uncertainty, and argument strength. *Journal of Experimental Psychology: Applied*, *15*, 199–212.

Durik, A.M., Britt, M.A., Reynolds, R., & Storey, J. (2008). The effect of hedges in persuasive arguments: A nuanced analysis of language. *Journal of Language and Social Psychology*, *27*, 217–234.

Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition*. Oxford, UK: Oxford University Press.

Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, *7*, 569–592.

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

Hume, D. (1977). *An enquiry concerning human understanding*. Indianapolis: Hackett. (Original work published 1748.)

Johnson, S.G.B., Jin, A., & Keil, F.C. (2014). Simplicity and goodness-of-fit in explanation: The case of intuitive curve-fitting. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Johnson, S.G.B., Johnston, A.M., Toig, A.E., & Keil, F.C. (2014). Explanatory scope informs causal strength inferences. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Johnson, S.G.B., Rajeev-Kumar, G., & Keil, F.C. (2014). Inferred evidence in latent scope explanations. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Johnson, S.G.B., Rajeev-Kumar, G., & Keil, F.C. (2015). *Sense-making under ignorance*. Manuscript under review.

Johnson, S.G.B., & Keil, F.C. (2015). *Evidential asymmetries in diagnostic reasoning*. Manuscript in preparation.

Johnston, A.M., Johnson, S.G.B., Koven, M.L., & Keil, F.C. (2015). Probabilistic versus heuristic accounts of explanation in children: Evidence from a latent scope bias. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, *116*, 20–58.

Khemlani, S.S., Sussman, A.B., & Oppenheimer, D.M. (2011). *Harry Potter* and the sorcerer's scope: Latent scope biases in explanatory reasoning. *Memory & Cognition*, *39*, 527–535.

Leslie, S. J. (2008). Generics: Cognition and acquisition. *Philosophical Review*, *117*, 1–47.

Lipton, P. (2004). *Inference to the best explanation* (2nd Edition.). London: Routledge.

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, *55*, 232–257.

McDonald, J., Samuels, M., & Rispoli, J. (1996). A hypothesis-assessment model of categorical argument strength. *Cognition*, *59*, 199–217.

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

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

Peirce, C.S. (1997). *Pragmatism as a principle and method of right thinking: The 1903 Harvard lectures on pragmatism*. Albany, NY: SUNY 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.