Cognitive Psychology 89 (2016) 39–70

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Sense-making under ignorance

Samuel G.B. Johnson *, Greeshma Rajeev-Kumar, Frank C. Keil

Department of Psychology, Yale University, United States

ARTICLE INFO

Article history:
Accepted 20 June 2016

Keywords:
Causal reasoning
Categorization
Explanation
Ignorance
Probabilistic reasoning

ABSTRACT

Much of cognition allows us to make sense of things by explaining observable evidence in terms of unobservable explanations, such as category memberships and hidden causes. Yet we must often make such explanatory inferences with incomplete evidence, where we are ignorant about some relevant facts or diagnostic features. In seven experiments, we studied how people make explanatory inferences under these uncertain conditions, testing the possibility that people attempt to *infer* the presence or absence of diagnostic evidence on the basis of other cues such as evidence base rates (even when these cues are normatively irrelevant) and then proceed to make explanatory inferences on the basis of the inferred evidence. Participants followed this strategy in both diagnostic causal reasoning (Experiments 1–4, 7) and in categorization (Experiments 5–6), leading to illusory inferences. Two processing predictions of this account were also confirmed, concerning participants' evidence-seeking behavior (Experiment 4) and their beliefs about the likely presence or absence of the evidence (Experiment 5). These findings reveal deep commonalities between superficially distinct forms of diagnostic reasoning—causal reasoning and classification—and point toward common inferential machinery across explanatory tasks.

1. Introduction

Across perception and cognition, we fill in details missing from our actual experience. In perception, we see illusory contours and infer continuities of forms; indeed, we fill in unattended

* Corresponding author at: 2 Hillhouse Ave., New Haven, CT 06520, United States.
E-mail address: samuel.johnson@yale.edu (S.G.B. Johnson).

http://dx.doi.org/10.1016/j.cogpsych.2016.06.004
0010-0285/

elements of our visual field so successfully that we fail to appreciate the sharp limits of our conscious awareness. Likewise, in cognition, we fill in narratives, scripts, and schemas almost continuously through our daily lives. Although these acts of filling in can create striking illusions and false memories (Loftus & Palmer, 1974; Simons & Levin, 1997), this filling in tendency is an essential tool for cognition: Sound strategies for inferring unknown information allow us to get by with limited information, while still effectively navigating the world.

Here, we argue that this sort of filling in strategy plays a key role in explanatory reasoning, guiding our inferences about causal explanations and likely categorizations of objects, with people reasoning about such explanations based on both the observed and *inferred* evidence. We show at the same time, however, ways in which this strategy can lead to error when people base these inferences on irrelevant information.

1.1. Sense-making under ignorance

We must often make sense of things in the face of incomplete evidence. For example, doctors diagnose diseases when some test results are unavailable or inconclusive, giving the diagnosis they believe most likely or prudent given the evidence at hand. Juries infer the most likely culprit on the basis of often-sketchy evidence, conflicting testimony, and lawyerly doubletalk. People debate about ultimate explanations (e.g., the existence of God or of multiple universes) in the face of these explanations' intrinsically unverifiable predictions (e.g., an afterlife or the splitting of universes). More mundanely but no less remarkably, we all infer other people's mental states on the basis of just a few clues, infer the categories of objects even when some features are indeterminate, and infer causes when some of their potential effects are unknown. Explanation with incomplete evidence is the norm in everyday cognition.

Consider a simple concrete example. Suppose two trial attorneys are presenting two competing theories of a case to the jury (see Fig. 1). If Professor Plum committed the crime (call this hypothesis H_N, because it makes a single, ***n***arrow prediction), then there would be a dent in the candlestick (call this evidence X). Alternatively, if Colonel Mustard committed the crime (hypothesis H_W because it makes two, ***w***ider predictions), then there would be a dent in the candlestick (X), as well as mud on the drawing room carpet (Z). The observations posited by each hypothesis are depicted in Fig. 1.

Clearly, if Plum and Mustard are the only potential culprits, then the key question is whether there was mud in the drawing room (Z), because only this evidence would distinguish between the two hypotheses. That is, learning about the dent in the candlestick (X) is not diagnostic, because this observation would be equally consistent with either hypothesis—learning that this effect was present would tend to confirm both hypotheses (equally) and learning that it was absent would tend to disconfirm both hypotheses (equally). But if we find out that the mud was present, this would be powerful evidence in favor of Mustard, and if we find out that the mud was absent, this would be powerful evidence in favor of Plum. More generally, we rely on diagnostic evidence for telling apart competing explanations, whether the explanations are unobservable mental states, object categories, or causal events.

Sometimes, however, this diagnostic evidence is unavailable. If the jury faces a situation in which the evidence unambiguously indicates a dented candlestick (X), but is silent on the issue of the mud (Z)—say, because the floor had been cleaned before the detectives thought to check for it—then the jury faces incomplete evidence. Here, normative probability theory tells us that we should think the explanations equally likely: If we had no reason to think Plum or Mustard was the more likely culprit before gathering evidence, then we still have no reason after learning about X, but remaining ignorant about Z.

However, human judgments do not always obey probability theory (e.g., Kahneman, Slovic, & Tversky, 1982). Instead, we often use simplifying heuristics that perform reasonably well under ecologically realistic conditions but are prone to error. In cases of incomplete evidence, people tend to choose explanations that do not imply unknown evidence (Khemlani, Sussman, & Oppenheimer, 2011; Sussman, Khemlani, & Oppenheimer, 2014)—that is, people think that Professor Plum is the most likely culprit in the above case, against the dictates of probability theory. This error—known as the *latent scope bias*—is surprising both because it seems to deviate so strikingly from normative