Acta Psychologica 177 (2017) 10–16

Contents lists available at ScienceDirect

Acta Psychologica

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

Not that neglected! Base rates influence related and unrelated judgments

Michał Białek

Centre for Economic Psychology and Decision Sciences, Kozminski University, Jagiellońska 59, 03-301 Warsaw, Poland

ARTICLE INFO

Keywords:
Judgment
Biases
Heuristics
Inverse fallacy
Base-rate neglect
Conditional probability

ABSTRACT

It is claimed that people are unable (or unwilling) to incorporate prior probabilities into posterior assessments, such as their estimation of the likelihood of a person with characteristics typical of an engineer actually being an engineer given that they are drawn from a sample including a very small number of engineers.

This paper shows that base rates are incorporated in classifications (Experiment 1) and, moreover, that base rates also affect unrelated judgments, such as how well a provided description of a person fits a stereotypical engineer (Experiment 2). Finally, Experiment 3 shows that individuals who make both types of assessments – though using base rates to the same extent in the former judgments – are able to decrease the extent to which they incorporate base rates in the latter judgments.

1. Introduction

Consider the following problem:

One thousand people were tested in a study. Participants were a mixture of engineers and lawyers. Jack is a randomly chosen participant in this study. Jack is 36 years old. He is not married and is somewhat introverted. He likes to spend his free time reading science fiction and writing computer programs. How likely is Jack to be an engineer?

The most plausible way to answer such a question it is to use one's knowledge about the characteristics of engineers and lawyers and to apply these stereotypes in classifying Jack as a member of one of these groups. Such classifications are probabilistic in nature, and usually noted *p*(engineer). Such an expression of probabilistic knowledge is simplified however, as every probability estimate is conditioned on a population (Caves, 1990). For example, people think about how likely Jack is to be an engineer given that he has particular characteristics (Gavanski & Hui, 1992). Thus, the probabilistic statement can be more accurately expressed as *p*(classification | population), i.e., *p*(engineer | people having these particular characteristics). The better that the description of Jack fits the stereotype of an engineer, the more likely one is to assume that Jack is an engineer. This is termed the representativeness heuristic (Kahneman & Tversky, 1972, 1973).

What is the "stereotype" on which the assessment is conditioned? By and large, the stereotype can be understood as the likelihood of a person having particular characteristics given that they are an example of a group, i.e., *p*(characteristics | engineer). In the engineer-lawyer problem a particular set of features, such as being an introvert with an interest in sci-fi, is highly probable given that the person described is drawn from a population of engineers, and thus fits the stereotype well. Judgments based on stereotypes and classifications seem to be inversely based on the same data. Moreover, people make use of stereotypes when asked to make classifications, but when asked about a stereotype they might use a particular example to construct it. In short, people think of a stereotypical engineer by recalling the characteristics of a person they are highly familiar with, and they classify a new person as an engineer if they fit this stereotype. The two probability assessments, namely *p*(characteristics | classification) and *p*(classification | characteristics) are inherently confounded. For example, people in Europe and the USA, where there have recently been terrorist attacks by Islamic extremists, might be unable to distinguish the conditional probability of a person being a terrorist given that they are a Muslim: *p*(classification | characteristics), and the probability of being a Muslim given that one is a terrorist: *p*(characteristics | classification). However, these probabilities are quite different, which can easily be illustrated by substituting "male" for "Muslim". Most terrorists are male while very few males are terrorists.

Researchers have been interested in how such classifications of a person as belonging to a particular class can take into account prior probabilities, with a base rate coming from a particular sample rather than from a sample of convenience. People are expected to use information provided by an experimenter, i.e., the introductory sentence in the engineer-lawyer problem is as follows: *One thousand people were tested in a study. Participants were 5 engineers and 995 lawyers.* Here, people are required to estimate *p*(engineer | sample). Doing so should lead them to respond that it is more likely that Jack is a lawyer (De Neys & Feremans, 2013; Pennycook & Thompson, 2016). The reasoning behind this is that - considering the overwhelming number of lawyers, it is more likely that a person is a non-stereotypical lawyer than a

E-mail address: mbialek@kozminski.edu.pl.

http://dx.doi.org/10.1016/j.actpsy.2017.04.003
Received 21 September 2016; Received in revised form 1 February 2017; Accepted 10 April 2017
Available online 18 April 2017

stereotypical engineer.

Research shows that people systematically fail to condition their probability estimates on provided base rates, thus displaying so called base rate neglect (Christensen-Szalanski & Bushyhead, 1981). Authors in the heuristics and biases tradition claim that social stereotypes are a vivid and easy to understand source of information and are therefore processed effortlessly, using so called Type 1 processes (Evans & Stanovich, 2013). In contrast, probability is an abstract concept and requires effort to understand, with so called Type 2 processes being used. This difference in difficulty of processing results in a preference for stereotypes over base rates (Barbey & Sloman, 2007).

Such claims have been systematically challenged over the years, since some data suggest that base rates are more likely incorporated in probabilistic judgments if an estimate conditioned on a sample which is relevant (Koehler, 1996; Laming, 2007). Some researchers have gone further, showing that this incorporation is effortless and intuitive (Pennycook & Thompson, 2012). For example, Pennycook, Trippas, Handley, and Thompson (2014) asked people to consider a set of base rate problems and either rely on (a) their beliefs or (b) statistics. Results showed that base rates affected posterior probability estimates under both instructions, suggesting that the processing of base rates has an intuitive character which does not require conscious deliberation (and therefore is a Type 1 rather than a Type 2 process), and interacts or competes with other intuitions for ultimate control over behaviour. Despite being available without effort, the salience of probabilistic intuitions of this type is low, and therefore they usually lose out in an internal conflict with stronger and more vivid intuitions such as beliefs or stereotypes (De Neys, 2014; De Neys & Białek, 2017; Pennycook, Fugelsang, & Koehler, 2015).

To sum-up this line of research, two sources of information compete over ultimate judgments (stereotypes and base rates). These differ in salience, but are claimed to be intuitive. The degree to which base rates have a major impact on judgments depends on many factors, such as their format of presentation (Gigerenzer & Hoffrage, 1995) sample representativeness (Obrecht & Chesney, 2013) and whether people are prompted to reflect (Obrecht & Chesney, 2016).

Such models describing base rate usage compare two probabilistic judgments*:* p(classification) and stereotype-fit, but neglect the fact that each probability is already conditioned on a particular sample. More specifically, a classification already has a base rate on which it is conditioned, although this base rate is subjectively selected by a person at the time the classification has to be made.

In the heuristics and biases tradition researchers require participants to use the abstract base rates they are presented with (e.g., the numbers of engineers and lawyers in a sample) and to draw conclusions based uniquely on these base rates. In other words, people should assess the probability of Jack being an engineer given that he is drawn from a particular, non-representative sample: p(engineer | sample in the experiment). Experiments show, however, that people condition their assessments on another more specific sample, i.e., p(engineer | sample with particular characteristics). Thus, base rate neglect does not reflect an inability to process probabilistic information, but an inability to substitute an experience-based base-rate with a newly presented base-rate.[1] Thus, such substitution is difficult for people, as is ignoring one's own beliefs when tracking logical validity in reasoning (Newstead, Pollard, Evans, & Allen, 1992; Trippas, Thompson, & Handley, 2016). Given this difficulty, the best a person can do is to combine the two sources of information to form one joint basis on which to condition their assessment. For example, a decision maker can evaluate how likely Jack is to be an engineer given his characteristics and that he comes from a sample underrepresented among engineers: p(engineer | characteristics, sample). So, for example, if, given a description, one thinks that the probability of a person being an engineer is 80%,[2] the probability estimate would be only slightly affected by the new base rate (i.e., 5/1000) rather than being an average of the two probabilities. The reason for this is that people update their probabilistic beliefs too slowly given new evidence (Oaksford & Chater, 2009), even when their prior beliefs are based merely on guessing (Krueger & Clement, 1994).

Thus, in base rate neglect tasks the new evidence (the base rate in the relevant sample) would be included in an internal model of the external world (stereotype) using a particular weight from a previous model and a weight for the new evidence. As in Bayesian statistical inference, where evidence updates a priori beliefs, individuals will update their a priori probability assessments in the light of a new piece of evidence (Dienes, 2011; Gigerenzer & Hoffrage, 1995). Alternatively, the updated beliefs might be a result of cognitive algebra (Anderson, 1974).[3] Here, stronger evidence would affect probability assessments to a greater degree than weaker evidence. Consistently automatic processing of base rates sometimes fails, for example, where a base rate is not extreme (e.g., 300/700) rather than being a typically extreme ratio (e.g., 5/995), see Pennycook, Fugelsang, and Koehler (2012). Also, the more related the sample base rate, or the less diagnostic the individuating source of information, the greater the reliance on base rates (Bar-Hillel, 1980; Koehler, 1996).

To summarize this introduction, research has consided three possibilities with respect to base rates: (1) people ignore base rates; (2) people have probability-driven intuitions which are low in salience; (3) people update their a priori beliefs in the light of new evidence, with different weights being assigned to different sources of information. The predictions of these models were investigated in the following three experiments.

2. Experiment 1

Given the above line of reasoning, it is to be expected that people make a particular probability assessment based on their beliefs, which they subsequently update to make particular classifications. For instance, people classify Jack as an engineer based on their belief about how representative an introvert sci-fi fan is as an example of an engineer. However, instead of this, in the base rate neglect tradition, people would be expected to use the base rate provided by the experimenter (e.g., 5 engineers and 995 lawyers), thus classifying Jack as a member of the more frequently represented group in the sample (i.e., lawyers). I speculated that individuals do not neglect the new evidence provided in instructions, but they consider the two base rates jointly. Consistent with this, I expected peoples' probability assessments to be updated bidirectionally according to the new evidence provided by the base rate in a sample (as shown previously by many resarchers, among them Bar-Hillel (1980); Fischhoff and Bar-Hillel (1984); Pennycook et al. (2014). Compared to the use of assumed a priori probabilities with no base rate information provided, a high base rate would be expected to increase, and a low base rate to decrease, the estimated probability of a case being an example of a particular group. The positive verification of this hypothesis would disprove Model 1 (the complete base rate neglect hypothesis), but could not distinguish between Models 2 and 3.

[1] There is an issue with problems where no previous base rate is available, such as the proportions of blue and green cabs in an unknown abstract city, as used in the well know cab problem of Gigerenzer and Hoffrage (1995). Research suggests, that in such cases people use a 50–50 base rate instead (Einhorn & Hogarth, 1985; Fischhoff & De Bruin, 1999).

[2] A plausible estimation, given that in pilot studies with similar material the description fitted a particular stereotype, with a rating of 8 points on a 1–10 scale, (De Neys & Glumicic, 2008).

[3] A discussion of the nature of the updating of beliefs is beyond the scope of this paper, as there are several – Bayesian and non-Bayesian – models that are under investigation (Douven & Schupbach, 2015a, 2015b; Douven & Wenmackers, 2015).