To be published in (2016) *Cognitive Illusions: Intriguing Phenomena in Thinking, Judgment, and Memory* (2nd ed.).,

Hove, UK: Psychology Press.

Ed., R. Pohl

Base-Rate Neglect

Gordon Pennycook[1] & Valerie A. Thompson[2]

[1]Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada
[2]Department of Psychology, University of Saskatchewan, Saskatoon, Saskatchewan, Canada

Corresponding author:
Gordon Pennycook
University of Waterloo - Psychology
200 University Avenue West
Waterloo, ON, N2L 3G1
gpennyco@uwaterloo.ca

The authors declare no conflict of interest.

Chapter 2

Base-rate neglect

Gordon Pennycook and Valerie A. Thompson

The “base-rate” refers to the *a-priori* probability of an event or outcome. For example, there are 19 professional hockey players who play for the Toronto Maple Leafs at any given moment during the hockey season. On game day, 38 out of 2.5 million people in Toronto are National Hockey League (NHL) players (i.e., the Leafs and their opponent). Thus, the base-rate probability that a randomly encountered person in Toronto on game day is a NHL player is 38/2,500,000 or .00152%. Base-rate neglect refers to the phenomenon whereby people ignore or undervalue that probability, typically in lieu of less informative, but more intuitively appealing information about an individual case (Kahneman & Tversky, 1973). Thus, even if a Toronto resident were to come across a tall, burly, hockey-stick wielding man wearing a Maple Leafs jersey, the probability that he actually plays for the team (and is not simply a fan wearing the jersey on his way to a recreational hockey game) is very small. An everyday example of how base-rates such as this can be neglected can be illustrated with a thought experiment.

Imagine owning a car that constantly breaks down and, after a few years of this, you have finally found someone to purchase it. This additional bit of money allows you to purchase a new car – one that will hopefully be more reliable – though you do not have a very large budget. You have narrowed the list of potential cars to two options (which are approximately the same cost): a Subaru and a Fiat. The most recent issue of *Consumer Reports* indicates that Subaru owners typically have fewer mechanical problems than do the Fiat owners and that the Subaru was more highly rated by experts. However, you also happen to have an uncle who once owned a Subaru.

He informs you that his Subaru had multiple very costly problems. His suggestion is to go with the Fiat, which he feels is a more reliable car.

Which car do you purchase?

There is a strong temptation in situations such as this to ignore or underweight the base-rate probability of mechanical issues (i.e., based on the large sample of owners' experiences and expert opinion described in *Consumer Reports*) in lieu of the more appealing single case (i.e., based on your trusted uncle's experience). Indeed, when given hypothetical scenarios of this sort, participants often choose the 'Fiat' response – that is, the car that is probabilistically more likely to have mechanical issues but that has an intuitive appeal (Fong, Krantz, & Nisbitt, 1986). Clearly, neglecting the base-rates can be expensive, if one opts for the repair-needy Fiat over the more reliable Subaru. The neglect or underweighting of base-rate probabilities has been demonstrated in a wide range of situations in both experimental and applied settings (Barbey & Sloman, 2007). In this chapter we will outline some of the ways that the base-rate fallacy has been investigated, discuss a debate about the extent of base-rate use, and, focusing on one particular form of base-rate neglect, we will outline recent work on the cognitive mechanisms that underlie the tendency to underweight or ignore base-rate information.

BASE-RATE NEGLECT IN MANY FORMS

The term base-rate neglect applies to any case where a prior probability is not sufficiently weighted in reasoning. As a consequence, base-rate neglect takes many forms, a selection of which is illustrated in Text box 2.1. The purpose of the first two problems (1-2) is to create a conflict between base-rate and individuating (stereotype) information and see what proportion of individuals select the base-rate response. The first example (1) is referred to as an "implicit base-rate" problem, because the relevant base-rate is not mentioned. Instead, there is an explicit description of a set of stereotypes (orderly, precise, etc.) that suggests to most people that "Person A" is more likely to be a statistics major. Although implicit, the base-rate is nonetheless relevant to deciding which option is more likely. At most universities, there are far more students in General Arts than there are in Statistics (the ratio at the University of Waterloo in Canada is ~24:1); this discrepancy is so large that an individual who is orderly, precise, etc. is far more likely to major in General Arts than Statistics, despite the stereotypical association with a Statistics major. Nonetheless, when students at the University of Waterloo were given a set of these problems, only 21% selected the response consistent with the base-rate (in this case, General Arts). Moreover, response time analyses indicated that participants did not appear to recognize the relevance of the base-rate probability (i.e., they spent the same amount of time reasoning as when the problem contained no conflict between base-rate and stereotype); indicating that the 21% of the time when base-rate responses were given likely resulted from individuals having atypical stereotypes and not an understanding of the base-rate (Pennycook, Fugelsang, & Koehler, 2012). Thus, at least about 80% of the students in the study completely neglected the base-rates, but probably more did as well.

Text box 2.1 Frequently investigated varieties of base-rate neglect

(1) Person 'A' was selected at random from a group consisting of all University of Waterloo students majoring in either GENERAL ARTS or STATISTICS. Person 'A' is orderly, organized, precise, practical and realistic. Is Person A's major more likely to be: GENERAL ARTS or STATISTICS? (Pennycook, Fugelsang, & Koehler, 2012)

(2) In a study 1000 people were tested. Among the participants there were 995 nurses and 5 doctors. Paul is a randomly chosen participant of this study. Paul is 34 years old. He lives in a beautiful home in a posh suburb. He is well spoken and very interested in politics. He invests a lot of time in his career. Is Paul more likely to be: a doctor or a nurse? (Pennycook & Thompson, 2012)

(3) The probability of breast cancer is 1% for a woman at age forty who participates in routine screening. If a woman has breast cancer, the probability is 80% that she will get a positive mammography. If a woman does not have breast cancer, the probability is 9.6% that she will also get a positive mammography. A woman in this age group had a positive mammography in a routine screening. What is the probability that she actually has breast cancer? _%. (Gigerenzer & Hoffrage, 1995)

Consider now the second example (2), a base-rate problem modelled on Kahneman and Tverksy's (1973) "lawyer-engineer problem" (described subsequently). This problem is very similar to the first example (1), with the key exception that the base-rate probability (995 nurses, 5 doctors) is explicitly stated. The conflict between base-rate probability (indicating that Paul is a nurse) and the stereotypical information (indicating that Paul is a doctor) is now, in theory, plainly obvious. Nonetheless, Pennycook et al. (2012) found that participants selected the base-rate response only 24% of the time on a set of problems like this. Thus, participants typically fail to sufficiently weight base-rate information even when the prior probability is extreme (the prior probability that Paul is a doctor is 0.5%) and explicitly stated in the problem. However, in this case, participants do take longer when giving the stereotypical response to versions of this problem where the base-rate and stereotypes point to alternative responses, suggesting that they do successfully recognize (at some level) that the base-rate conflicts with the stereotype (Pennycook et al., 2012).

The third example (3) – referred to as the "mammography problem" (e.g., Eddy, 1982) – is quite different in form and plainly more complex than the previous two. The problem contains a base-rate (1% of women have breast cancer), but also includes information about the hit-rate (80% chance of a positive mammogram if breast cancer is present) and the false-alarm rate (9.6% chance of a positive mammogram if breast cancer is absent). Given the hypothesis (H) that a random 42 year old woman has a positive mammogram (the observed datum, D), the probability that she actually has breast cancer [P (H │D)] can be determined using Bayes' theorem (the specific details of which are not important for present purposes; interested readers can see Birnbaum, 2004; Kurzenhäuser & Lücking, 2004): (0.01)(0.80) / [(0.01)(0.80) + (0.99)(0.096)]. Based on this calculation, there is a 7.8% chance that the woman has breast cancer given a positive mammogram. Given the complexity of this operation, it is perhaps no surprise that very few people are able to generate the correct solution (e.g., 16% in Gigerenzer & Hoffrage, 1995). Indeed, even physicians have a great deal of difficulty with problems such as this (Hammerton, 1973). The median response on these types of problem is to report a number close to the hit-rate (i.e., 80%; Barbey & Sloman, 2007), ignoring the fact that the base rate indicates that the probability of cancer is very rare.

Text box 2.2 A classroom demonstration of base-rate neglect based on Kahneman & Tversky (1973)

Method

Participants

This is a between-participant experiment and requires three roughly equal groups. Fortunately, the phenomenon under investigation is very robust and groups of 10 or more individuals should suffice. If necessary, the experiment can be changed to a within-participant design and participants can complete each condition in order (Condition 1, Condition 2, Condition 3).

Materials and Procedure

Participants in each of the conditions will be given a slightly different task. It would be best for each participant to only see the task assigned to their condition, although the experiment should work regardless. The conditions are as follows:

Condition 1

Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feeling and little sympathy for other people and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense.

How similar is Tom W. to the typical graduate student in [your country] in each of the following nine fields of graduate specialization? Please rank the following nine fields of graduate specialization in order of the relative similarity of Tom W. relative to the prototypical student in [your country]. Rank from 1 to 9, using each rank only once.

Business Administration
Computer Science
Engineering
Humanities and Education
Law
Library Science
Medicine
Physical and Life Sciences
Social Science and Social Work

Condition 2
Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feeling and little sympathy for other people and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense.

The preceding personality sketch of Tom W. was written during Tom's senior year in high school by a psychologist, on the basis of projective tests. Tom W. is currently a graduate student in [your country]. Please rank the following nine fields of graduate specialization in order of the likelihood that Tom W. is now a graduate student in each of these fields. Rank from 1 to 9, using each rank only once

Business Administration
Computer Science
Engineering
Humanities and Education
Law
Library Science
Medicine
Physical and Life Sciences
Social Science and Social Work

Condition 3
Consider all first year graduate students in [your country] today. Please write down your best guesses about the percentage of these students who are now enrolled in each of the following nine fields of specialization:

Business Administration
Computer Science

Engineering
Humanities and Education
Law
Library Science
Medicine
Physical and Life Sciences
Social Science and Social Work

Note: If universities in your country do not offer exactly these fields of specialization, please replace them with the fields that come closest.

Analysis

Participants in Condition 3 provided subjective base-rates. Note that it does not matter if these perceived base-rates are accurate; they just need to accurately represent the typical opinions of the participants in the other conditions. This can be double-checked by having participants in Conditions 1 (the “similarity” group) and 2 (the “likelihood” group) also complete Condition 3. To compute the base-rates, one needs to compute a mean of the estimated base-rates for each graduate specialization (in %). To compute the measure of similarity and of likelihood, compute the mean ranks for each graduate specialization for Conditions 1 and 2 (respectively).

On this task, base-rate neglect occurs when participants’ likelihood ratings are informed by the similarity of each person to a stereotype rather than the base-rates. To demonstrate this, one needs to correlate the likelihood judgments with both the similarity rankings and the base-rate estimates. For this, put the responses for each condition in separate columns of the same table (see Table 2.1). If participants used the base-rate of graduate specialization (i.e., the responses for Condition 3) when judging the likelihood of graduate specialization (Condition 2), then there should be a positive correlation between the responses for Conditions 2 and 3. If, on the other hand, participants used stereotypes (i.e., the responses for Condition 1) to determine the likelihood of graduate specialization, there should be a positive correlation between the responses for Conditions 1 and 2. Finally, these two correlation coefficients can be compared to assess which source of information was more influential. If the similarity was more influential than the base-rate probability, the correlation between responses for Conditions 1 and 2 should be larger than the correlation between responses for Conditions 2 and 3.

Results and Discussion

This experiment was the first in Kahneman and Tversky’s (1973) seminal work on base-rate neglect. Their results can be found in Table 2.1. They found a very strong positive correlation between similarity (Condition 1) and likelihood (Condition 2), $r = .97$. The rankings in the two groups were nearly identical! In contrast, not only was there no *positive* correlation between mean judged base-rate (Condition 3) and likelihood judgments (Condition 2), but the correlation was actually *negative*, r = -.65. This is because Tom W. sounds most like someone in computer science or engineering (associated with relatively low base-rates) and least like someone in humanities and social sciences (associated with relatively high base-rates). Clearly, base-rates were not taken into account when participants were asked to judge the *likelihood* that Tom W. was a student in these graduate specializations. Thus, this is an example of base-rate neglect.

Table 2.1. Estimated base-rates of the nine areas of graduate specialization and summary of similarity and likelihood ratings for Tom W. (Kahneman & Tversky, 1973).

Graduate specialization area	Mean similarity rank	Mean likelihood rank	Mean judged base-rate (in %)
Business Administration	3.9	4.3	15
Computer Science	2.1	2.5	7
Engineering	2.9	2.6	9
Humanities and Education	7.2	7.6	20
Law	5.9	5.2	9
Library Science	4.2	4.7	3
Medicine	5.9	5.8	8
Physical and Life Sciences	4.5	4.3	12
Social Science and Social Work	8.2	8.0	17

THEORETICAL ACCOUNTS

The term “base-rate neglect” implies that information about base-rates is completely ignored. In this section, we will summarize research that examines whether this is true or whether the term “neglect” is a bit of a misnomer. We will then move to some more recent research on the cognitive mechanisms that underlie base-rate neglect. As was made evident in Text box 2.1, there are many different manifestations of base-rate neglect. Not surprisingly, therefore, there is no unified theoretical account of base-rate neglect. Instead, the theorizing tends to be centred on the cognitive mechanisms thought to underlie particular forms of base-rate neglect, which may not be applicable to other forms.

Are base-rates ignored?

The earliest data on base-rate neglect seemed to indicate that people essentially ignore base-rates when making judgments. The Tom W. problem (Text box 2.2) from Kahneman and Tversky (1973) is a particularly striking example of this. In the case of the more complex mammography problem (Text box 2.1) (and others like it), Eddy (1982) found that fewer than 5% of respondents were able to correctly solve the problem and Hammerton (1973) found only nominally better performance in a group of physicians (10% correct). Results such as this led many researchers to conclude that base-rates are ignored. However, subsequent research showed that this conclusion was too pessimistic (e.g., Gigerenzer & Hoffrage, 1995).

For example, there are a number of conditions under which people can be made sensitive to base-rate information (Birnbaum, 2004): Participants are more sensitive to base-rates when given multiple problems with varying base-rate probabilities (Fischhoff & Bar-Hillel, 1984) or if given problems where the base-rates come *after* individuating information (e.g., stereotypes; Krosnick, Li, & Lehman, 1990). Sensitivity to base-rates is also facilitated by manipulations that make a causal link between the base-rate and the judged case explicit (e.g., Bar-Hillel, 1980). Consider the two examples in Text box 2.3. In the first example (1), the Cab problem, the color distribution of the cabs is the base-rate information (85% blue, 15% green) and the accuracy of witness identification is the individuating information (80% hit-rate and 20% false-alarm rate). According to Bayes' theorem, the probability that the cab was green is 41% because the base-rate and individuating information needs to be integrated [(0.8/0.2) * (0.15/0.85)]. However, the typical response for this problem is 80% (i.e., the hit rate; Bar-Hillel, 1980). However, if there is a causal link between the base-rate and individuating information, people are more inclined to combine them. To understand why, consider the second example (2) in Text box 2.3, the Motor problem. This problem is exactly the same in terms of base-rate (A = 85%, B = 15%) and individuating information (80% hit-rate and 20% false-alarm rate), and the correct answer is therefore also 41%. The key difference between the problems is that the Motor problem makes it clear that the base-rate is an *attribute* of the two motors. Or, in other words, it is readily apparent that the base-rate is causally linked to the function of the motors. This manipulation highlighted the importance of the base-rate and, as a consequence, over 60% of the participants' gave a response that indicated sensitivity to the base-rate (Bar-Hillel, 1980).

Text box 2.3 Causality in base-rate problems (Bar-Hillel, 1980)

(1) Two cab companies operate in a given city, the Blue and the Green (according to the color of cab they run). 85% percent of the cabs in the city are Blue, and the remaining 15% are Green. A cab was involved in a hit-and-run accident at night. A witness later identified the cab as a Green cab. The court tested the witness' ability to distinguish between Blue and Green cabs under night time visibility conditions. It found that the witness was able to identify each color correctly about 80% of the time, but confused it with the other color about 20% of the time.
What do you think are the chances that the errant cab was indeed Green, as the witness claimed?

(2) A large water-pumping facility is operated simultaneously by two giant motors. The motors are virtually identical (in terms of model, age, etc.), except that a long history of breakdowns in the facility has shown that one motor, call it A, was responsible for 85% of the breakdowns, whereas the other, B, caused 15% of the breakdowns only. To mend a motor, it must be idled and taken apart, an expensive and drawn out affair. Therefore, several tests are usually done to get some prior notion of which motor to tackle. One of these tests employs a mechanical device which operates, roughly, by pointing at the motor whose magnetic field is weaker. In 4 cases out of 5, a faulty motor creates a weaker field, but in 1 case out of 5, this effect may be accidentally caused. Suppose a breakdown has just occurred. The device is pointed at motor B.
What do you think are the chances that motor B is responsible for this breakdown?

These results would not be expected if base-rates are completely ignored. Importantly, there is also evidence that base-rates can not only enter into judgment, but that people are capable of using them correctly. Consider a modified version of the mammography problem:

> 10 out of every 1,000 women at age forty who participate in routine screening have breast cancer. 8 out of every 10 women with breast cancer will get a positive mammography. 95 out of every 990 women without breast cancer will also get a positive mammography. Here is a new representative sample of women at age forty who got a positive mammography in routine screening. How many of these women do you expect to actually have breast cancer? __ out of __

Here, the probabilities (1%; 80%; 9.6%) have been presented in terms of natural frequency formats (10 out of 1,000; 8 out of 10; 95 out of 990). This relatively straightforward manipulation was sufficient to increase performance by a factor of ~3 (46% v. 16% accuracy; Gigerenzer & Hoffrage, 1995, see also Kurzenhäuser & Lücking, 2004; Tversky & Kahneman, 1983). On the basis of these results, Gigerenzer & Hoffrage (1995) argued that frequency formats are more easily understood by participants because they are consistent with the sequential way that information is acquired in the context of natural sampling. This view is typically associated with evolutionary psychology and, in particular, the idea that humans have evolved an intuitive way of dealing with base-rates that requires the right sort of conditions to be triggered but that does not require conscious deliberation (see Barbey & Sloman, 2007, for a review).

Now consider Kahneman and Tverksy's (1973) "lawyer-engineer" problem (mentioned above):

> A panel of psychologists have interviewed and administered personality tests to 30 engineers and 70 lawyers, all successful in their respective fields. On the basis of this information, thumbnail descriptions of the 30 engineers and 70 lawyers have been written.

> You will find on your forms a description, chosen at random from the 100 available descriptions. Please indicate your probability that the person described is an engineer, on a scale from 0 to 100.
>
> The same task has been performed by a panel of experts, who were highly accurate in assigning probabilities to the various descriptions. You will be paid a bonus to the extent that your estimate comes close to those of the expert panel.
>
> Here is the description: Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing, and mathematical puzzles. The probability that Jack is one of the 30 engineers in the sample of 100 is ___%.

This problem contains base-rate information that is easily understood, such that most participants successfully give the base-rate rate response on versions of this problem that do not contain stereotypes (Pennycook & Thompson, 2012). Nonetheless, as discussed above, participants do not typically give the base-rate response when base-rates are in conflict with stereotypical information. Rather, they rely primarily on the representative information and underweight (but not necessarily neglect) the base-rate.

In a recent experiment, Pennycook and Thompson (2012) investigated the degree to which participants used base-rates in a set of problems (18 in total) of the lawyer-engineer type. They included two between-participant conditions (see Text box 2.4): A standard base-rate (BR) condition (1) and a no base-rate (NoBR) condition (3). The goal of this manipulation was to see what sort of influence base-rates had on probability judgments. If base-rates are completely ignored, judgments should not differ between conditions.

Text box 2.4 Conditions from Pennycook and Thompson (2012)

(1) In a study 1000 people were tested. Among the participants there were 995 nurses and 5 doctors. Paul is a randomly chosen participant of this study. Paul is 34 years old. He lives in a beautiful home in a posh suburb. He is well spoken and very interested in

politics. He invests a lot of time in his career. What is the probability (0-100) that Paul is a nurse? [Base-rate Condition (BR); Incongruent]

(2) In a study 1000 people were tested. Among the participants there were 995 who live in a condo and 5 who live in a farmhouse. Kurt is a randomly chosen participant of this study. Kurt works on Wall Street and is single. He works long hours and wears Armani suits to work. He likes wearing sunglasses. What is the probability (0-100) that Kurt lives in a condo? [Base-rate Condition (BR); Congruent]

(3) In a study 1000 people were tested. Among the participants there were nurses and doctors. Paul is a randomly chosen participant of this study. Paul is 34 years old. He lives in a beautiful home in a posh suburb. He is well spoken and very interested in politics. He invests a lot of time in his career. What is the probability (0-100) that Paul is a nurse? [No base-rate Condition (NoBR)]

In addition, Pennycook and Thompson (2012) also included a within-subject manipulation of congruency such that in the BR condition the base-rates were sometimes inconsistent with the stereotypes, akin to the lawyer-engineer problem, and sometimes consistent with the stereotypes. For example, in Text box 2.4 the first problem (1) is incongruent because the base-rate indicates that Paul is very likely to be a nurse but the stereotypes suggest that Paul sounds more like a doctor. In contrast, the second problem (2) is congruent because the base-rate indicates that Kurt is likely to own a condo and the stereotypes are more consistent with a condo owner than a farmhouse owner. [Note: Congruency could not be manipulated in the NoBR condition due to the lack of base-rates.] This manipulation was included because it is possible that base-rates may be used differently depending on their association with the individuating stereotypical information.

Pennycook and Thompson's (2012) key results can be found in Figures 2.1 and 2.2., which show the distribution of probability estimates across the conditions. The comparison between the distributions of probability estimates given base-rates (BR) or not (NoBR) shows that base-rates had a substantial influence on judgments. Moreover, the form that this took differed depending on whether base-rates and stereotypes were consistent (congruent) or inconsistent (incongruent). As is evident from Figure 2.1, the vast majority of probability estimates for congruent problems in the BR condition were 90% or higher (*Mean* = 88.6%). In contrast, probability estimates ranged fairly equally from 50-90% when participants were only given stereotypical information in the NoBR condition (*Mean* = 68.5%). This pattern of results indicates that base-rates not only informed participants' judgments, but that the modal response was a *combination* of base-rate and individuating (stereotypical) information. Participants integrated the two sources of information, as is necessary for Bayes' theorem.

Figure 2.1. Distribution of probability estimates for congruent problems (from "Reasoning with base rates is routine, relatively effortless, and context dependent" by Gordon Pennycook and Valerie A. Thompson, 2012, *Psychonomic Bulletin & Review*, *19*, 531, © Psychonomic Society, Inc. 2012. Adapted with permission of the publisher). For BR condition, high responses are consistent with both stereotypes and base-rates. For No-BR condition, high responses are consistent with stereotypes. Note that problems were not "congruent" in the No-BR condition due to the lack of base-rate information. They are the exact congruent problems from the base-rate condition with base-rates removed. The counterbalancing was such that, in the no base-rate condition, the problems would have been "congruent" or "incongruent" if base-rates had been included.

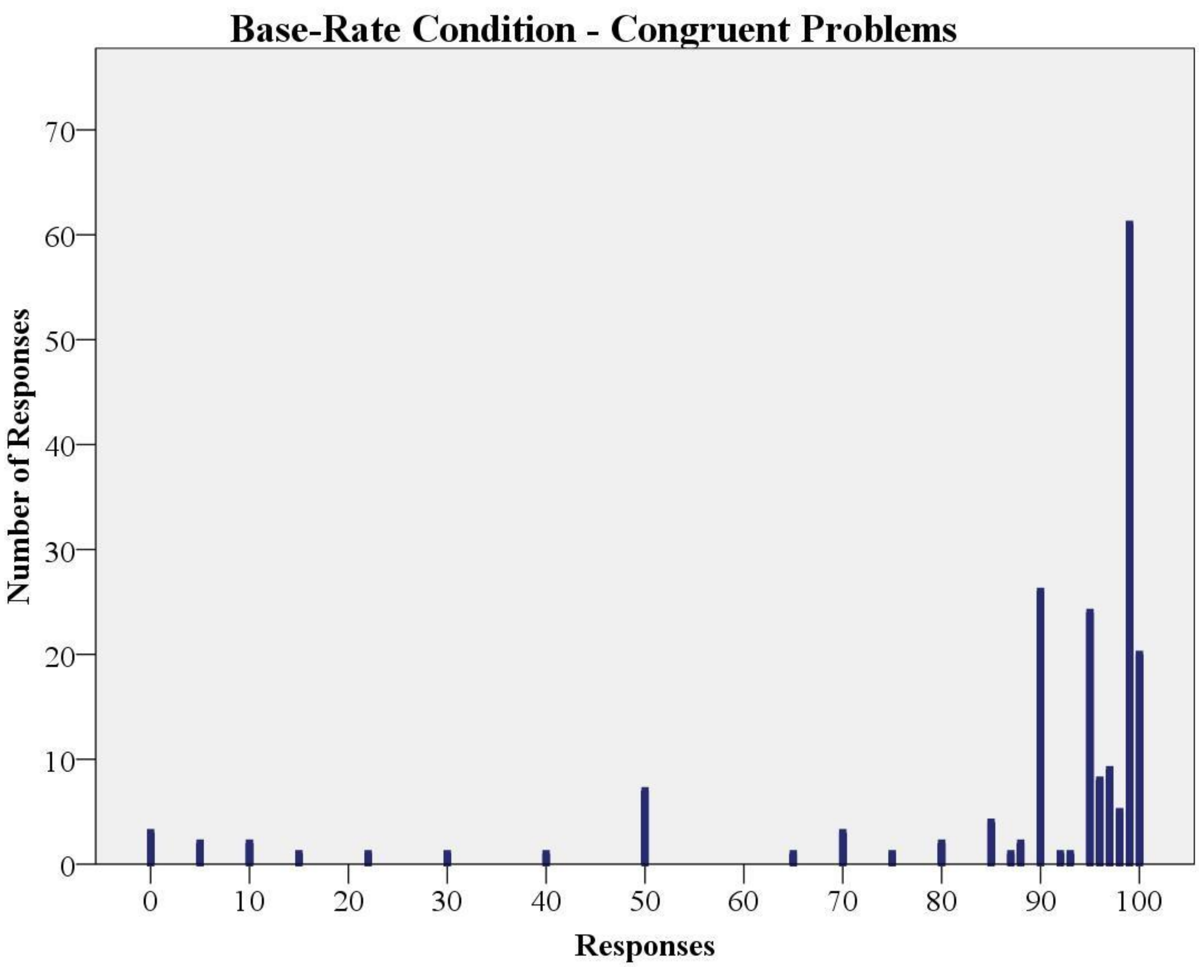

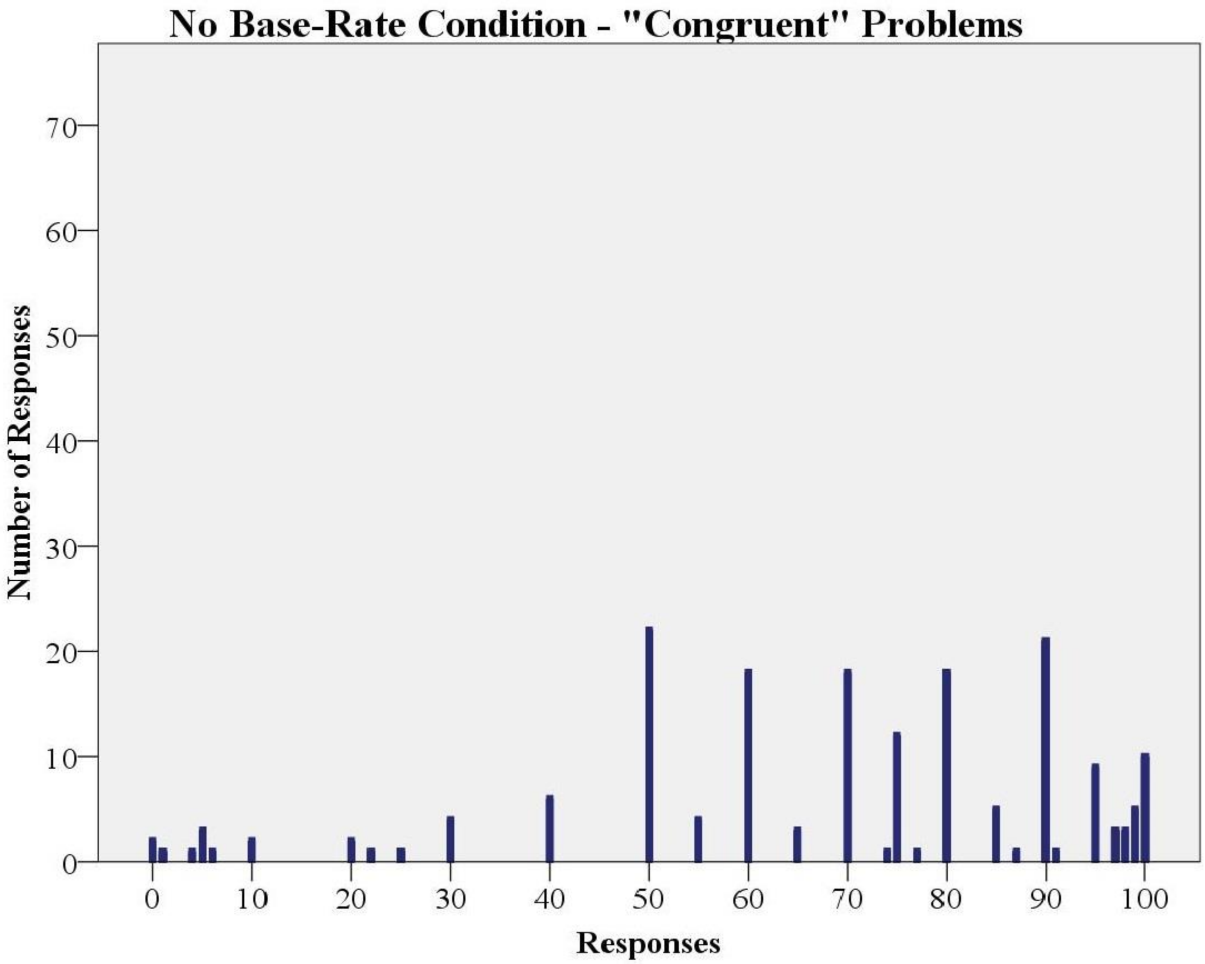
No Base-Rate Condition - "Congruent" Problems
Number of Responses
0
10
20
30
40
50
60
70
0
10
20
30
40
50
60
70
80
90
100
Responses

An entirely different pattern was evident when base-rates and stereotypes were inconsistent (incongruent). In the BR condition (Figure 2.2.), one cluster of responses was similar to that found for the congruent problems (i.e., responses 90% and higher); however, there was a second cluster more consistent with traditional base-rate neglect findings. Namely, many responses were 10% and lower, which is very *inconsistent* with the presented base-rates (and therefore very *consistent* with the presented stereotypes). This contrasts starkly with the responses for the No-BR conditions and the congruent BR condition. It seems that when base-rates and stereotypes point to different responses, participants no longer integrate them. Rather, they select one or the other source of information and give a relatively extreme response. If participants successfully integrated base-rate and individuating information for incongruent problems most of the probability judgments would be clustered somewhere in the middle of the distribution. As is evident from Figure 2.2., this did not happen. This pattern of results indicates that base-rates are sometimes very influential (i.e., when they are consistent with stereotypes) but are also sometimes neglected in lieu of stereotypes. Moreover, some people give base-rate responses even when base-rates and stereotypes are in conflict. This may be due, in part, to the fact that participants received multiple problems with slightly variable base-rate information (Kahneman & Frederick, 2005). Thus, to understand these findings, we need to move past the debate about whether base-rates are ignored and into a more theoretical discussion of the cognitive mechanisms that underlie the use (or neglect) of base-rates.

Figure 2.2. Distribution of probability estimates for incongruent problems (from "Reasoning with base rates is routine, relatively effortless, and context dependent" by Gordon Pennycook and Valerie A. Thompson, 2012, *Psychonomic Bulletin & Review*, *19*, 532, © Psychonomic Society, Inc. 2012. Adapted with permission of the publisher). For the BR condition, high estimates are consistent with base-rates and low estimates are consistent with stereotypes. For the No-BR condition, high responses are consistent with stereotypes. Note that problems were not "incongruent" in the No-BR condition due to the lack of base-rate information. They are the exact incongruent problems from the base-rate condition with base-rates removed. The counterbalancing was such that, in the no base-rate condition, the problems would have been "congruent" or "incongruent" if base-rates had been included.

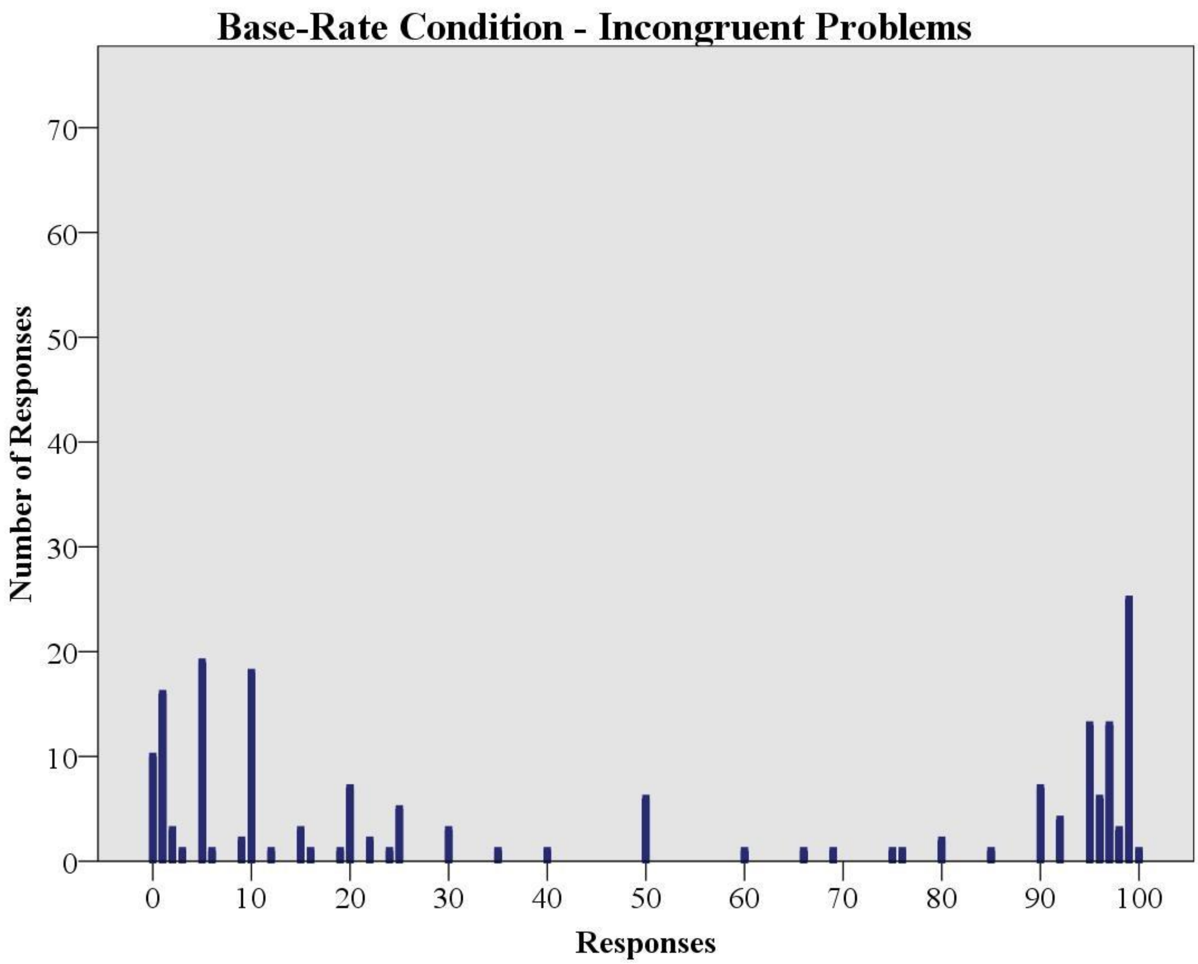

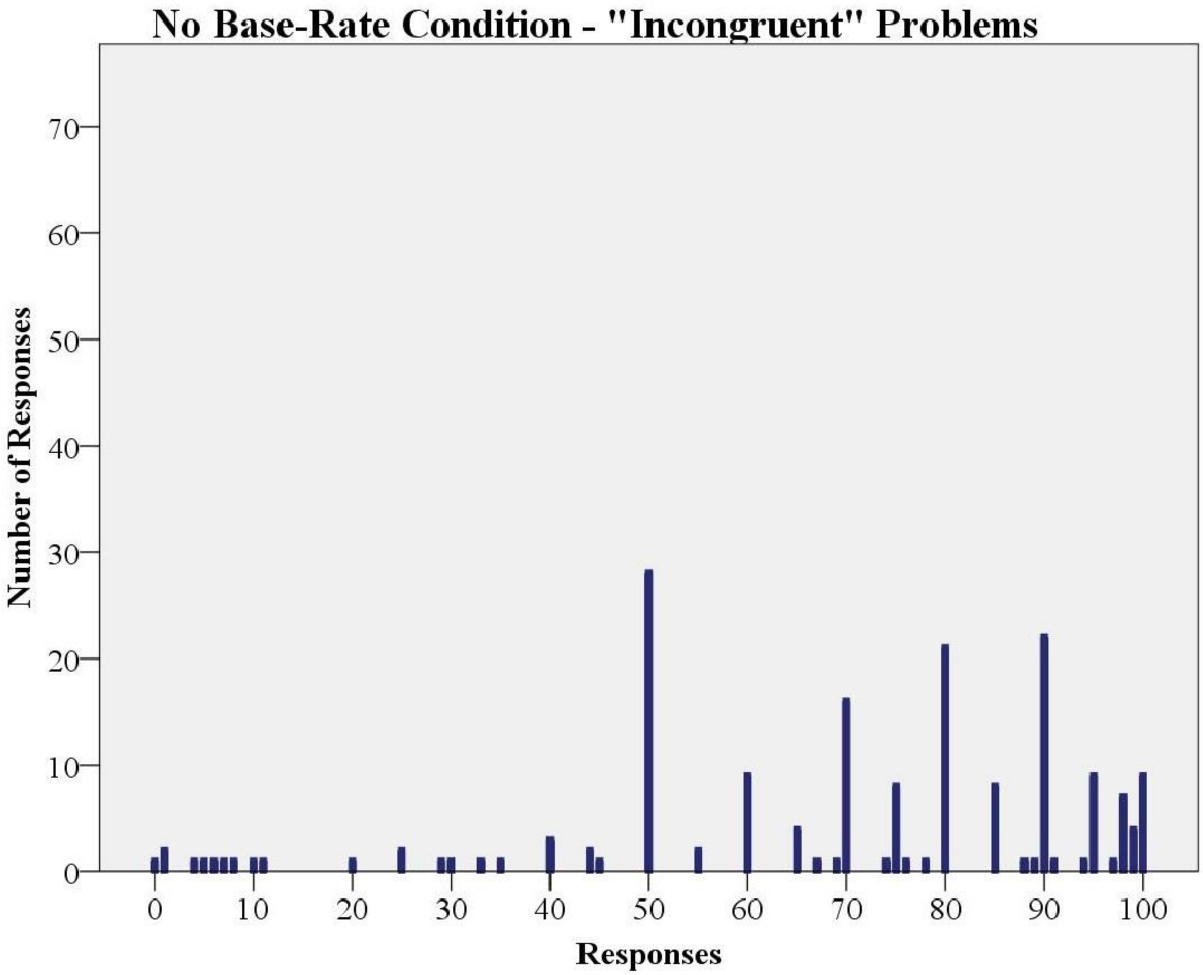
No Base-Rate Condition - "Incongruent" Problems
Number of Responses
70
60
50
40
30
20
10
0
0
10
20
30
40
50
60
70
80
90
100
Responses

Dual-process theory and base-rate neglect

The dominant explanation for why some types of individuating information (e.g., stereotypes) are favored so heavily over base-rates appeals to dual-processing. Dual-process theory relates to the idea that there are two types of processes by which humans make judgments and decisions (Evans & Stanovich, 2013): Type 1 processes that are autonomous, fast, and high capacity, and Type 2 processes that are reflective, slow, and resource demanding. The role of Type 1 processes are to provide default outputs which can be accepted, rejected, or modified as explicit representations in working memory via Type 2 processing (Evans & Stanovich, 2013).

The initial explanation of base-rate neglect anticipated these later developments in dual-process theory. Namely, participants were thought to form a rapid response using a "representativeness heuristic" (Kahneman & Tversky, 1973; see Chapter 11 in this volume). That is, rather than answering the rather difficult question regarding the probability of group membership, participants formed their judgment on the basis of which group the personality description seemed more representative. This explanation is quite amendable with dual-process theory. Specifically, stereotypes cue an intuitive "Type 1" response (based on representativeness) and base-rates require deliberate "Type 2" reasoning processes to enter into judgment (e.g., Kahneman, 2003). Since humans typically forego costly Type 2 processing in favor of less effortful Type 1 processing (Stanovich & West, 2000), stereotypes are naturally favored over base-rates. Indeed, participants who are more disposed to analytic thought (as indexed by both self-report and performance measures) are more likely to give the base-rate response for problems of the lawyer-engineer type (e.g., Pennycook, Cheyne, Barr, Koehler, & Fugelsang, 2014).

Although all seem to agree that individuating information like stereotypes are very intuitive sources of information, there is clearly some disagreement about how difficult base-rates are to use. Kahneman's (2003) dual-process account holds that base-rates require resource demanding reasoning processes whereas other accounts hold that base-rates do not require any deliberation at all and are actually quite intuitive (at least, when they are in the right format; e.g., Gigerenzer & Hoffrage, 1995). Fortunately, recent experiments have started to clarify this issue.

Recall the Pennycook and Thompson (2012) experiment where participants were given problems with base-rates (BR) or without base-rates (NoBR; Text box 2.4). The researchers also included an additional within-participant manipulation that was quite revealing. Participants were asked to respond to each problem twice: First they provided whatever response initially popped into their head (an intuitive response given under a time deadline) and then they responded to the same question again with a final answer given over free time. When offered this chance to rethink their intuitive response, participants were just as likely to shift toward the stereotype as they were toward the base-rate. Moreover, many participants gave responses consistent with the base-rates even when they gave the first response that came to mind. These results indicate that responses based on both stereotypes and base-rates can be either intuitive *or* reflective.

This conclusion was supported by a set of experiments by Pennycook, Trippas, Handley, and Thompson (2014). Participants were given a set of base-rate problems of the lawyer-engineer type and were explicitly instructed how to respond to each problem: 1) Statistics instructions highlighted the importance of base-rates in determining the likelihood of group membership, and 2) Belief instructions highlighted the importance of belief-based information (stereotypes) in determining the likelihood of group membership. If base-rates require slow Type

2 processing and belief judgments are made using fast Type 1 processing, then responding according to the base-rates should not interfere with judgments when responding based on beliefs. Instead, across three experiments, participants had just as much difficulty responding according to belief instructions as they did with statistics instructions. Namely, probability estimates were less accurate, confidence was lower, and response time was longer when base-rates and stereotypes conflicted *regardless of the instruction manipulation*. This result was replicated when participants were put under a strict time deadline. This represents rather striking evidence that the use of base-rates can be intuitive.

If base-rates use is (at least sometimes) intuitive, how can we explain the preponderance of stereotypical responses to problems of the lawyer-engineer type? The answer to this question requires a more nuanced understanding of what it means for something to be "intuitive" (Thompson, 2014). It may be that responses based on *both* stereotypes and base-rates can be intuitive, but that the former are typically *more* accessible in that stereotypes cue a response that comes to mind more quickly and fluently than the response cued by the base-rate information (Pennycook, Fugelsang, & Koehler, 2015). This would leave the stereotype as the default response and, as a consequence, even if participants recognized the importance of the base-rates they would still need to inhibit and override the default stereotypical response (De Neys & Franssens, 2009). Since humans are miserly information processors (Stanovich & West, 2000), this resource demanding Type 2 response is often foregone; hence base-rate neglect.

The miserly processing account explains why base-rate responses are more common among individuals who are more disposed to analytic thought (e.g., Pennycook, Cheyne, et al., 2014) and more intelligent (e.g., Thompson & Johnson, 2014). It also explains why base-rate responding has been linked with additional psychological factors. For example, those who are less prone to base-rate neglect are more likely to be skeptical of religious and paranormal claims (e.g., Pennycook, Cheyne, et al., 2014). In other words, those who are more likely to question their initial intuitions about stereotypes in the context of a base-rate problem are also more likely to question widely held and often quite intuitive supernatural beliefs. Moreover, people who are better able to detect the conflict between base-rates and stereotypes may also be better able to detect the intrinsic conflict between ubiquitous materialistic intuitions (e.g., that beings cannot pass through solid objects) and immaterial beliefs (e.g., that an angel can pass through solid objects; Pennycook, Cheyne, et al., 2014). The low-level conflict between base-rates and stereotypes evident for (at least some) base-rate problems may also be a key trigger of Type 2 processing (i.e., something that *causes* people to think; Pennycook et al., 2015) and the way people respond to this type of conflict (which would presumably be prevalent in many different domains, such as the conflict between supernatural and materialist beliefs) is a key aspect of human cognition. Moreover, an understanding of these mechanisms could be used to potentially devise interventions that help reduce or even overcome base-rate neglect. For example, in order to facilitate conflict detection, Pennycook et al. (2015) gave participants multiple problems with extreme base-rates that were presented *after* stereotypes. In that condition, participants actually gave *more* base-rate responses than stereotypical ones.

CONCLUSIONS

Base-rate neglect is a very robust phenomenon that comes in many forms (Barbey & Sloman, 2007). Nonetheless, much evidence suggests that base-rates are not always neglected (e.g., Gigerenzer & Hoffrage, 1995). Although the calculation of probability estimates may require deliberative reasoning at least some of the time for at least some people, base-rates influence judgment at a lower, more intuitive level (Pennycook, Trippas, et al., 2014). Further, people appear to be able to detect the conflict between base-rates and individuating information (e.g., stereotypes) that is common in certain forms of base-rate neglect (Pennycook et al., 2012). This conflict detection is an important bottom-up source of analytic reasoning (Pennycook et al., 2015) and, as a consequence, base-rate neglect has been linked to psychological phenomena not typically associated with reasoning and decision making (e.g., religious belief; Pennycook, Cheyne et al., 2014). Thus not only does base-rate neglect have important consequences in applied areas (e.g., medical decision making; Eddy, 1982), but at a more theoretical level, the study of base rate neglect has revealed novel insights about the interaction between deliberate and analytic thinking in ways that has informed our understanding of a wide array of cognitive illusions and reasoning biases.

SUMMARY

- People often neglect or underweight base-rate probabilities when other (typically more intuitive) information is available.
- Base-rate neglect has been demonstrated using a wide range of tasks across many experiments.
- There are ways to improve people's reasoning with base-rates, though they are typically still underweighted.
- Base-rates can be processed without deliberative reasoning (though typically not as intuitively as stereotypical individuating information).
- Base-rate neglect emerges as a consequence of an interaction between intuitive and reflective processes.

FURTHER READING

Barbey and Sloman (2007) represents an extensive review of base-rate neglect research in the context of competing models. Pennycook, Trippas, Handley, and Thompson (2014) outline typical dual-process account of base-rate neglect and provide evidence for a revised version of that model. Pennycook, Fugelsang, and Koehler (2015) use base-rate problems to illustrate the key role of conflict detection as a source of analytic engagement.

ACKNOWLEDGEMENTS

The cited studies of the authors were funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada. This came in the form of a master’s and doctoral funding for GP and a Discovery Grant for VT.

REFERENCES

Barbey, A. K., & Sloman, S. A. (2007). Base-rate respect: From ecological rationality to dual processes. *Behavioural and Brain Sciences, 30*, 241-256.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica, 44*, 211-233.

Birnbaum, M. H. (2004). Base rates in bayesian inference. In R. F. Pohl (Ed.), *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory* (pp. 43-60). New York, NY: Psychology Press.

De Neys, W., & Franssens, S. (2009). Belief inhibition during thinking: Not always winning but at least taking part. *Cognition, 113*, 45–61.

Eddy, D. M. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, & A. Tversky (Eds), *Judgment under uncertainty: Heuristics and biases* (pp. 249–67). New York: Cambridge University Press.

Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives in Psychological Science,* 8, 223–241.

Fischhoff, B., & Bar-Hillel, M. (1984). Focusing techniques: A shortcut to improving probability judgments? *Organizational Behavior and Human Performance, 34*, 175-194.

Fong, G. T., Krantz, D. H., & Nisbitt, R. E. (1986). The effects of statistical training on thinking about everyday problems. *Cognitive Psychology, 18*, 253-292.

Gigerenzer, G. & Hoffrage, U. (1995) How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review, 102*, 684–704.

Hammerton, M. (1973). A case of radical probability estimation. *Journal of Experimental Psychology, 101*, 252–54.

Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, *58,* 697-720.

Kahneman, D. & Frederick, S. (2005). A model of heuristic judgement. In K. J. Holyoak & R. G. Morrison (Eds.), *The Cambridge Handbook of Thinking and Reasoning* (pp. 267–293). Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 237–251.

Krosnick, J. A., Li, F., & Lehman, D. R. (1990). Conversational conventions, order of information acquisition, and the effect of base rates and individuating information on social judgments. *Journal of Personality and Social Psychology, 59*, 1140-1152.

Kurzenhäuser, S., & Lücking, A. (2004). Statistical formats in Bayesian inference. In R. F. Pohl (Ed.), *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory* (pp. 61–77). New York, NY: Psychology Press.; in Pohl R (ed): Cognitive

Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., & Fugelsang, J.A. (2014). Cognitive style and religiosity: The role of conflict detection. *Memory & Cognition, 42,* 1-10.

Pennycook, G., Fugelsang, J. A., & Koehler, D. J. (2012). Are we good at detecting conflict during reasoning? *Cognition, 124,* 101–106.

Pennycook, G., Fugelsang, J. A., & Koehler, D. J. (2015). What makes us think? A three-stage dual-process model of analytic engagement. *Cognitive Psychology, 80*, 34-72.

Pennycook, G., & Thompson, V. A. (2012). Reasoning with base rates is routine, relatively effortless, and context-dependent. *Psychonomic Bulletin and Review, 19*, 528–534.

Pennycook, G., Trippas, D., Handley, S. J., & Thompson, V. A. (2014). Base-rates: Both neglected and intuitive. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 40*, 544-554.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences, 23,* 645-726.

Thompson, V. (2014). What intuitions are… and are not. In Brian H. Ross (Ed.), *The Psychology of Learning and Motivation* (pp. 35-75). Burlington: Academic Press.

Thompson, V., & Johnson, S. C. (2014). Conflict, metacognition, and analytic thinking. *Thinking & Reasoning*, *20*, 215-244.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 91,* 293-315.