What is base rate neglect example

Statistical formal fallacy The fallacy of the base rate, also called the neglect of the base rate or the distortion of the base rate, is a fallacy. When related base rate information (i.e. general prevalence information) and specific information (i.e. information related only to a specific case) are presented, people tend to ignore the base rate in favor of individuating the information, rather than properly integrating the two. [1] The neglect of the basic rate is a specific form of more general neglect of enlargement. False positive paradox An example of the fallacy of the base rate is a false positive paradox. This paradox describes situations where there are more false positive test results than actual positive results. For example, 50 out of 1,000 people have a positive test for infection, but only 10 have an infection, meaning that 40 tests were false positives. The probability of a positive test result is determined not only by the accuracy of the test, but also by the characteristics of the sample set. [2] If prevalence, the proportion of those who have the condition is lower than the false positive test rate, even tests that have a very low chance of false positive in an individual case give overall more false than actual positive results. [3] The paradox surprises most people. [4] It is particularly counterintuitive in interpreting a positive test result on a population with low prevalence after looking at positive results from a population with high prevalence. [3] If the false positive test rate is higher than the proportion of the new population with the condition, the test administrator whose experience has been drawn from testing in a high prevalence population may infer from experience that a positive test result is usually indicated by a positive subject, when in fact it is much more likely that a false positive subject has occurred. Examples Example 1: Disease High-incidence population Number of people Infected Uninfected Total Testpositive 400 (true positive) 30 (false positive) 430 Testnegative 0(false negative) 570 (true negative) 570 Total 400 600 1000 Imagine that running an infectious disease test on population A 1000 people in which 40% of the test is infected has a false positive rate of 5% (0.05) and no false negative rate. The expected result of 1000 tests for population A would be: Infected and the test shows the disease (true positive) 1000 × 40/100 = 400 people would receive a truly positive Uninfected and the test suggests (false positive) 1000 × 100 – 40/100 × 0.05 = 30 people would receive a false positive The remaining 570 tests are correctly negative. So, in population A, a person who receives a positive test can be more than 93% confident (400/30 + 400) that it correctly indicates infection. Low population Number of people infected with uninfected total positive test 20 (true positive) 49 (false positive) 69 Test negative 0(false negative) 931 (true 931 A total of 20,980,1000 now consider the same test applied to population B, in which only 2% are infected. The expected result of 1000 tests for population B would be: Infected and the test shows the disease (true positive) 1000 × 2/100 = 20 people would receive a truly positive Uninfected and the test indicates a disease (false positive) 100 0 × 100 - 2/100 × 0.05 = 49 people would receive false positive Tests the remaining 931 (= 1000 - (49 + 20)) are correctly negative. In population B, only 20 out of a total of 69 people with a positive test result are actually infected. So the probability that they actually become infected after a person is told that a person is infected is only 29% (20/20 + 49) for a test that otherwise seems to be 95% accurate. A tester with group A experience might consider it a paradox that in Group B, the result that usually correctly indicated infection is usually false positive. Confusion of the posterior probability of infection with the previous probability of receiving a false positive is a natural mistake after receiving the result of a health-threatening test. Example 2: Drunk drivers A group of police officers have breathing apparatus showing false drunkenness in 5% of cases where the driver is sober. However, breathalyzers never fail to detect a really drunk person. One in 1,000 drivers drive drunk. Suppose the officers then stop the driver randomly to perform a breath test. That means the driver's drunk. We assume you don't know anything else about them. How high is the probability that they're really drunk? Many would answer up to 95%, but the correct probability is about 2%. The explanation is as follows: on average, for every 1000 drivers tested, 1 driver is drunk and it is 100% certain that for this driver the actual positive test result is, so there is 1 actual positive test result of 999 drivers not drunk, and among those drivers there are 5% false positive test results, so there are 49.95 false positive test results Therefore, the probability that one of the drivers between 1 + 49.95 = 50.95 positive test results really is drunk is 1 / 50.95 ≈ 0.019627 {\displaystyle 1/50.95\approx. 0.019627} . However, the validity of this result depends on the validity of the original assumption that the policeman actually stopped the driver accidentally, and not because of a bad ride. If this or any other involuntable reason for stopping the driver was present, then the calculation also includes the probability that the drunk driver drives competently and the busy driver drives (in-) competently. More formally, the same probability of approximately 0.02 can be determined using bayes' sentence. The goal is to determine the probability that the driver is drunk due to the fact that breathalyzer stated that they are drunk, which can be represented as p (d you n k ∣ D) {\displaystyle p(\mathrm {drunk} \mid where D means that the breathing apparatus indicates that the driver is drunk. Bayes's theorm tells us that p (d you are n k ∣ D) = p (P (D ∣ d you n k) p (d you n k) p (D) . {\displaystyle p(\mathrm {drunk} \mid D)={\frac {p(D\mid \mathrm {drunk})\,p(\mathrm {drunk})}{p(D)}}.} In the first paragraph we were told the following: p (d are you n k) = 0.001 , {\displaystyle p(\mathrm {drunk})=0.001,} p (s o b e r) = 0.999 , {\displaystyle p(\mathrm {sober })=0.999,} p (D ∣ d are you n k) = 1.00 , {\displaystyle p(D\mid \mathrm {drunk})=1.00,} and p (D ∣ s o b e r) = 0.05. {\displaystyle p(D\mid \mathrm {sober})=0.05.} As you can see from the formula, one needs a p(D) for bayes theorem, which can be calculated from the previous values using the Total Probability Act: p (D) = p (D ∣ d are you n k) p (d are you n k) + p (D ∣ s o b e r) p (s o b e r) {\displaystyle p(D)=p(D mid\\mathrm drunk {} ,p(\mathrm {drunk})+p(D\mid \mathrm {sober})\,p(\mathrm {sober})} giving p (D) = (1.00 × 0.001) + (0.05 × 0.999) = 0.05095. {\displaystyle p(D)=(1.00\times 0.001)+(0.05\times 0.999)=0.05095.} By attaching these numbers to the Bayes theorm, one finds that p (d you are n to ∣ D) = 1.00 × 0.001 0.05095 = 0.019627. {\displaystyle p(\mathrm {drunk} \mid D)={\frac {1.00\times 0.001}{0.05095}}=0.019627.} Example 3: Terrorist identification In a city of 1 million people, there were 100 terrorists and 999,900 were not terrorists. To simplify the example, it is assumed that all the people present in the city are residents. Thus, the probability that a randomly selected city resident is a terrorist is 0.0001, and the probability of the base rate that the same population is not a terrorist is 0.9999. In an effort to catch terrorists, the city is installing an alarm system with a security camera and automatic facial recognition software. The software has two failure rates of 1%: False negative rate: If the camera scans a terrorist, it rings 99% of the time and does not ring 1% of the time. False positive rate: If the camera scans no terrorist, the bell does not ring 99% of the time, but rings 1% of the time. Let's assume the residents set off the alarm now. What are the chances that this person is a terrorist? In other words, what is P(T ∣ B), the probability that the terrorist was detected due to the ringing of the bell? Anyone making a 'basic rate of deception' would infer that there is a 99% chance that the identified person is a terrorist. While the conclusion seems to make sense, it's actually bad reasoning, and the calculation below will show that the chances of being terrorists are actually close to 1%, not nearly 99%. The fallacy stems from the confusing nature of two different failures. The number of non-ringing per 100 terrorists and the number of non-terrorists per 100 bells are unrelated One does not necessarily equal the other, nor does it have to be almost the same. To show this, consider what would happen if the same alarm system was set up in the second city, without terrorists. As in the first city, alarm sounds for 1 in every 100 non-terrorist residents detected, but unlike the first city, the alarm never sounds for terrorists. Therefore, 100% of all cases where there is an alarm sound, for not terrorists, but false negative speed can not even be calculated. The number of non-terrorists per 100 bells in this city is 100, but P (T ∣ B) = 0 %. There's zero chance the terrorist was exposed, given the bell ringing. Imagine that the first city's entire population of one million people pass in front of the camera. About 99 out of 100 terrorists will raise the alarm – and so there will be about 9,999,900 non-terrorists. Therefore, about 10,098 people will raise the alarm, among which about 99 will be terrorists. So the probability that the person who sets off the alarm is actually a terrorist is only about 99 out of 10,098, which is less than 1% and very, very much below our original estimate of 99%. The deception of the basic rate is so misleading in this example, because there are far more terrorists than terrorists, and the number of false positives (not terrorists being scanned as terrorists) is much greater than the actual positives (the actual number of terrorists). Findings in Psychology In experiments, it was found that people prefer individualating information over general information when the first is available. [5] [6] [7] In some experiments, students were asked to estimate the averages of points (GPAs) of hypothetical students. With relevant GPA distribution statistics, students tended to ignore them if they received descriptive information about a particular student, even though the new descriptive information was clearly of little or no relevance to the school's performance. [6] This finding was used to claim that interviews are an unnecessary part of the college admissions process because interviewers are unable to select successful candidates better than basic statistics. Psychologists Daniel Kahneman and Amos Tversky tried to explain this finding in the sense of a simple rule or heuristic called representativeness. They argued that many judgments on probability or cause and effect are based on how representative one thing or category is. [6] Kahneman considers the failure of the basic rate to be a specific form of neglect of enlargement. [8] Richard Nisbett argued that some attributions to prejudice, such as the basic error of attribution, are cases of basic rate deception: people do not use consensual information (base rate) about how others behaved in similar situations and instead prefer simpler dispositional attributions. [9] There is considerable debate in psychology conditions under which people value or do not value basic rate information. [10] [11] In the Heuristics and Prejudice Program, researchers highlighted empirical findings that show that people tend to ignore base rates and draw conclusions that violate certain norms of probate reasoning, such as the Bayes theorem. The conclusion drawn from this line of research shone a light on the fact that human probability thinking is fundamentally flawed and prone to errors. [12] Other researchers emphasized the link between cognitive processes and information formats, arguing that such conclusions are not generally justified. [13] [14] Consider example 2 again from above. The desired conclusion is to estimate the (rear) probability that the (randomly selected) driver is drunk, given that the breath test is positive. Formally, this probability can be calculated using the Bayes theorm, as mentioned above. However, there are different ways to present the relevant information. Consider the following, formally equivalent variant of the problem: 1 in 1,000 drivers drive drunk. Breathing apparatus can never detect a really drunk person. For 50 of the 999 drivers who are not drunk, breathalyzer falsely shows drunkenness. Suppose the police officers then stop the driver randomly, and force them to take a breath test. That suggests they're drunk. We assume you don't know anything else about them. How high is the probability that they're really drunk? In this case, the relevant numerical information — p(drunk), p(D ∣ drunk), p(D ∣ sober) — is presented in terms of natural frequencies with respect to a particular reference class (see reference class problem). Empirical studies show that people's conclusions correspond more closely to the Bayes rule when information is presented in this way, which helps to overcome base rate neglect in laymen[14] and professionals. [15] As a result, organizations like Cochrane Collaboration recommend using this kind of format to communicate health statistics. [16] Teaching people to translate these kinds of Bayesian reasoning problems into natural frequency formats is more effective than simply learning to involve probabilities (or percentages) in bayesian theorms. [17] It has also been shown that graphical representation of natural frequencies (e.g. icon fields) helps people to derive better. [17] [18] [19] Why are natural frequency formats useful? One important reason is that this information format facilitates the desired conclusion by simplifying the necessary calculations. This can be seen when using an alternative method of calculating the required probability p (drunk| D): p (d are you n k ∣ D) = N (d are you n k ∩ D) N (D) = 1 51 = 0.0196 {\displaystyle p(\mathr {drunk} \mid D)={\\frac {N(\mathrm {drunk} \cap D)}{N(D)}}={\frac {1}{51}}=0.0196}, where N(drunk ∩ D) indicates the number of drunk drivers and the result of breathalyzer and N(D) indicates the total number of cases with a positive breathalyzer result. The equivalence of this equation with the above results from the theory of probability axioms, according to which N(drunk ∩ D) = N × p (D ∣ drunk) × p (drunk). Importantly, while this equation is formally equivalent to the Bayes rule, it is not psychologically equivalent. The use of natural frequencies simplifies derivation because the required mathematical operation can be performed on natural numbers instead of normalized fractions (i.e. probabilities), because this makes a high number of false positives transparent and because natural frequencies show a nested structure. [20] [21] Not every frequency format facilitates Bayesian thinking. [21] [22] Natural frequencies refer to information on the frequency resulting from natural sampling[23] which stores information on the basic rate (e.g. number of drunk drivers taking a random sample of drivers). This differs from systematic sampling, at which the basic rates are set a priori (e.g. in scientific experiments). In the latter case, it is not possible to derive a posterior probability p (drunk ∣ positive test) from a comparison of the number of drunk and positive tests compared to the total number of persons who obtain a positive breath stimulus result, since the information on the basic rate is not retained and must be explicitly reintroduced using the Bayes theorm. See also Bayesian Probabilities Bayes' Theorem Data Dredging Induction Argument List of Cognitive Biases List of Paradoxes Misguided Vibrancy Prevention Paradox Plaintiff's Delusion, a reasoning error that involves ignoring the low previous probability of the Simpson Paradox, another error in statistical reasoning dealing with comparison groups Stereotype Links ^ Logical Fallacy: Basic Rate Deception. Fallacyfiles.org. Won 2013-06-15. ↑ Rheinfurth, M. H.; Howell, L.W. (March 1998). Probability and statistics in aerospace engineering (PDF). Nasa. p. 16. REPORT: False positive tests are more likely than actual positive tests if the overall population has a low prevalence of the disease. It's called a false positive paradox. ^ a b Vacher, H. L. (May 2003). Quantitative literacy - drug testing, cancer screening and identification of screed rocks. Journal of Geoscience Education: 2. At first glance, this seems perverse: the fewer students as a whole use steroids, the more likely a student identified as a user will not be a user. This was called a false positive paradox - Quote: Gonick, L.; Smith, W. (1993). Caricature guide statistics. Harper Collins. p. 49. ↑ Madison, B. L. (August 2007). Mathematical knowledge of citizenship. In Schoenfeld, A. H. (ed.). Evaluation of mathematical expertise. Publications of the Research Institute of Mathematical Sciences (New ed.). Cambridge Press. p. 122. ISBN 978-0-521-69766-8. The correct [probability estimate...] is surprising to many; therefore, the term paradox. ^ Bar-Hillel, Maya (1980). The fallacy of the base rate in probate courts (PDF). Acta Psychologica. 44 (3): 211-233. doi:10.1016/0001-6918(80)90046-3. ↑ a b c Kahneman, Daniel; Amos Tversky (1973). The psychology of prediction. Psychological overview. 80 (4): 237-251. doi:10.1037/h0034747. S2CID 17786757. ↑ Kahneman, Daniel; Amos Tversky (1985). The evidentiary impact of base rates. In Daniel Kahneman, Paul Slovic & Amos Tversky (ed.). Judgment in uncertainty: Heuristism and prejudice. Science. 185. pp. 153-160. doi:10.1126/science.185.4157.1124. PMID 17835457. S2CID 143452957. ↑ Kahneman, Daniel (2000). Rating by moment, past and future. In Daniel Kahneman and Amos Tversky (ed.). Options, values, and frames. ^ Nisbett, Richard E.; E. Borgia; R. Crandall; H. Žad (1976). Popular induction: Information is not always informative. In J.S. Carroll & J. W. Payne (ed.). Knowledge and social behavior. 2. pp. 227-236. ↑ Koehler, J.J. (2010). Intoxication of the base rate reassessed: descriptive, normative and methodological challenges. Behavioral and brain science. 19: 1-17. doi:10.1017/S0140525X00041157. S2CID 53343238. ↑ Barbey, A. K.; Sloman, S. A. (2007). Basic rate of respect: from ecological rationality to dual processes. Behavioral and brain science. 30 (3): 241-254, debate 255-297. doi:10.1017/S0140525X07001653. PMID 17963533. S2CID 31741077. ↑ Tversky, A.; Kahneman, D. (1974). Judgment in uncertainty: heuristism and prejudice. Science. 185 (4157): 1124–1131. Bibcode:1974Sci... 185.1124T. doi:10.1126/science.185.4157.1124. PMID 17835457. S2CID 143452957. ↑ Cosmides, Leda; John Tooby (1996). Are people good intuitive statisticians after all? Rethinking some of the conclusions of the literature on judgment in uncertainty. Knowledge. 58: 1–73. CiteSeerX 10.1.1.131.8290. doi:10.1016/0010-0277(95)00664-8. S2CID 18631755. ↑ a b Gigerenzer, G.; Hoffrage, USA (1995). How to improve Bayesian reasoning without instructions: Frequency formats. Psychological overview. 102 (4): 684. CiteSeerX 10.1.1.128.3201. doi:10.1037/0033-295X.102.4.684. ^ Hoffrage, United States of America; Lindsey, S.; Hertwig, R.; Gigerenzer, G. (2000). Medicine: Disclosure of statistical information. Science. 290 (5500): 2261-2262. doi:10.1126/science.290.5500.2261. PMID 11188724. S2CID 33050943. ↑ Akl, E. A.; Oxman, A. D.; Herrin, J.; Vist, G. E.; Terrenato, I.; Sperati, F.; Costiniuk, C.; Blank space, D.; Schünemann, H. (2011). Schünemann, Holger (šp.). Use alternative statistical formats to present risks and reduce risk. Cochrane Database of Systematic Evaluations (3): CD006776. doi:10.1002/14651858.CD006776.pub2. PMC 6464912. PMID 21412897. ↑ a b Sedlmeier, P.; Gigerenzer, G.

(2001). Teaching Bayesian Reasoning in Less two hours. Journal of Experimental Psychology: General. 130 (3): 380. doi:10.1037/0096-3445.130.3.380. hdl:11858/00-001M-0000-0025-9504-E. ^ Brase, G. L. (2009). Pictorial representation in statistical reasoning. Applied cognitive psychology. 23 (3): 369-381. doi:10.1002/acp.1460. S2CID 18817707. ↑ Edwards, A.; Elwyn, G.; Mulley, A. (2002). Risk explanation: Turn numeric data into meaningful images. BMJ. 324 (7341): 827-830. doi:10.1136/bmj.324.7341.827. PMC 1122766. PMID 11934777. ↑ Girotto, V.; Gonzalez, M. (2001). Solving probate and statistical problems: the question of the information structure and form of questions. Knowledge. 78 (3): 247-276. doi:10.1016/S0010-0277(00)00133-5. PMID 11124351. S2CID 8588451. ↑ a b Hoffrage, U.; Gigerenzer, G.; Krauss, S.; Martignon, L. (2002). Representation makes it easier to think: What natural frequencies are and what they are not. Knowledge. 84 (3): 343-352. doi:10.1016/S0010-0277(02)00050-1. PMID 12044739. S2CID 9595672. ↑ Gigerenzer, G.; Hoffrage, USA (1999). Overcoming difficulties in Bayesian reasoning: the answer of Lewis and Keren (1999) and Mellers and McGraw (1999). Psychological overview. 106 (2): 425. doi:10.1037/0033-295X.106.2.425. hdl:11858/00-001M-0000-0025-9CB4-8. ^ Kleiter, G. D. (1994). Natural sampling: Rationality without basic rates. Contributions to mathematical psychology, psychometry and methodology. Recent research in psychology. 375-388. doi:10.1007/978-1-4612-4308-3_27. ISBN 978-0-387-94169-1. External Links Fallacy Base Rate Deception Files Retrieved From