Transcripts of Training Videos

Table of Contents

Expert Fallibility Training

Hi there, I'm John Wilcox, and I'm a cognitive psychologist at Stanford University.

In this video, I'm going to talk about the fallibility of opinions from experts—such as scientists, academics, doctors and others.

Experts are arguably important. We sometimes use their opinions to inform us about things we are otherwise less knowledgeable about. To that extent, arguably society needs experts.

However, sometimes experts have fallible opinions.

For example, in one study, 118 US physicians were asked to make diagnoses of medical conditions [(Meyer et al., 2013)]. Some of these were about hard cases where approximately 5% of their diagnoses were correct, and others were about easier cases where approximately 55% of the diagnoses were correct. The concerning finding was that these doctors were similarly confident in the hard cases where they were likely to be wrong as they were in the easier cases where they were likely to be right. So not only did they make inaccurate misdiagnoses, but they were similarly confident in the hard cases where they were likely to make such misdiagnoses. On the basis of this and other findings, the study authors claimed that the physicians were generally "characterized by overconfidence in [their] accuracy". [(Meyer et al., 2013, p. 1952)] This shows that medical experts, including doctors, can have fallible opinions and overconfidence in their judgments.

Likewise, another study investigated the accuracy of hundreds of experts about political topics. Over a 20-year period, these political experts made predictions about events, such as the outcomes of elections or wars. Some of the experts were very inaccurate when making long-term predictions. For example, when they were 100% confident that something would not happen, that thing actually did happen 19% of the time. So they were overconfident. Furthermore, experts were not necessarily more accurate just because they had traditional indicators of expertise. As the study author, Philip Tetlock notes:

> It made virtually no difference whether participants had doctorates, whether they were economists, political scientists, journalists, or historians, whether they had policy experience or access to classified information, or whether they had logged many or few years of experience in their chosen line of work. [(Tetlock, 2005, p.68)]

And these are not the only studies showing that experts can have overconfident and inaccurate judgments. For example, another study examined the accuracy of 1,629 anesthesiologists from 80 countries, and it reported that 92% of them were overconfident [(Naguib et al., 2019)]. Another study examined the accuracy of physicians' judgments about cardiopulmonary variables [(Perel et al., 2016)]. The physicians were again inaccurate, and the study authors claimed this demonstrates the "very limited clinical ability of physicians to correctly assess important physiological parameters" [(Perel et al., 2016, p. 517)]. Berner & Graber (2008) also conducted a review of the medical literature, and they concluded that "physicians in general underappreciate the likelihood that their diagnoses are wrong" (p. 52).

So here are the main takeaway messages about experts. Sometimes experts have inaccurate or overconfident opinions. This is true in medicine when we consider, for example, doctors. It is also true in politics when we consider, for example, political commentators. And it is also true in other areas.

Additionally, non-experts sometimes have more accurate opinions. For example, during the 2020 pandemic, Youyang Gu developed a model to predict COVID deaths. It turned out that this model was more accurate than many other models that were developed by medical and epidemiological experts. For example, his model predicted that the US would record 231,000 deaths on the first of November, 2020. This prediction was highly accurate, as the US reported 230,995 deaths on the first of November, 2020. The difference between his prediction and the actual number was only 5 deaths. However, Youyang Gu was then 26-years old, with no qualifications in epidemiology or medicine. So he was not considered an expert in that sense. Instead, he had a background in data science. This, then, is a case where a non-expert had more accurate predictions than many experts in the field.

To summarize, expertise is often a valuable and necessary feature of society. But despite this, sometimes experts have mistaken opinions or overconfidence about topics in their area of expertise, and sometimes non-experts have trustworthy opinions in fields which they are not experts in.

Thank you for watching this video. Please answer the questions below this video and carry on with the study.

Reasoning Training

Hi there, I'm John Wilcox, and I'm a cognitive psychologist at Stanford University.

In this video, I'm going to talk about base rate neglect. Base rate neglect is a so-called fallacy—that is, an error in our reasoning. The purpose of this video is to teach you the fallacy, so that you can avoid it. To do this, we are going to walk you through the following hypothetical scenario.

Imagine that Sally takes a test for breast cancer. If she has breast cancer, there is 90% probability she would test positive. Suppose Sally then tests positive. What is the probability that Sally has breast cancer? An intuitive answer is to say she has a 90% probability of breast cancer. But this intuitive answer is incorrect. It results from base rate neglect. The correct answer is actually that we simply cannot determine the probability that Sally has breast cancer given the positive test. We're going to explain why this is so. To do this, imagine the following is true.

Suppose that out of every 1000 women, only 1 percent or 10 of them have breast cancer. This means that 99 percent or 990 of them do not have breast cancer. This statistic is called the *base rate* of breast cancer among women. The base rate for a disease is the proportion of people in a population who have the disease. Remember this!

Anyway, as mentioned before, imagine that if a woman has breast cancer, then there is a 90% probability that she would test positive. This means that for every 10 women who have breast cancer, 9 would test positive. So let us imagine that 9 of the 10 women with breast cancer would test positive.

Also imagine that if a woman does not have breast cancer, then she would test positive with a probability of 10%. If this is the case and 990 out of 1000 women do not have breast cancer, then 10 percent or 99 of those 990 would test positive. So let us imagine that 99 out of the 990 women without cancer would test positive. That means 99 women out of every 1000 would test positive even though they do not have breast cancer.

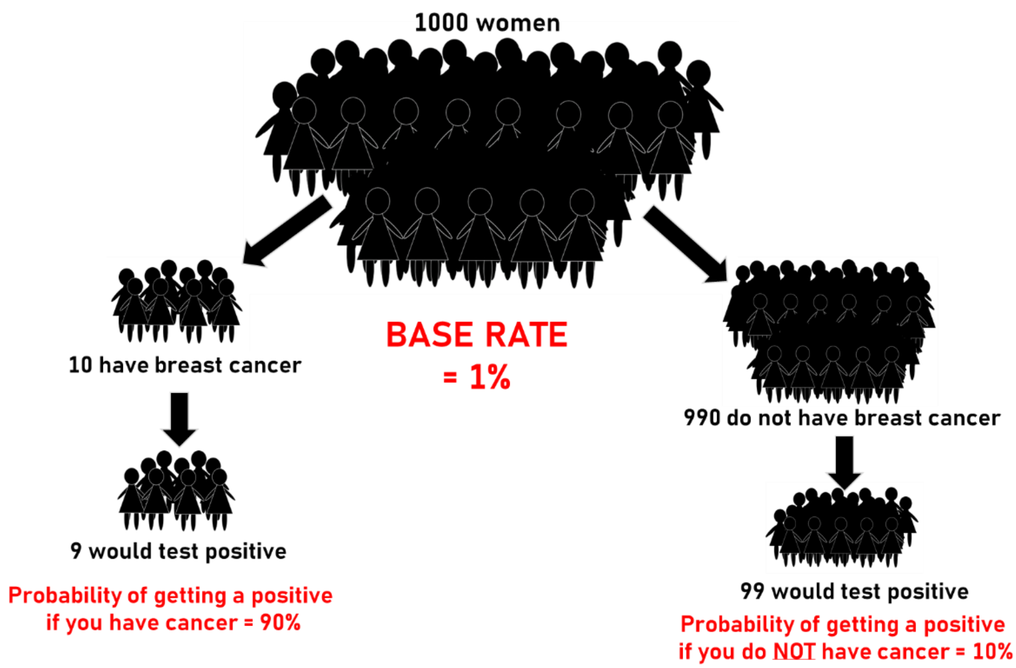

If a woman tests positive, they might be one of the 9 women who would test positive and do have breast cancer, but it is actually much more likely that they would be one of the 99 women who do not have breast cancer but would still test positive. For that reason, if a woman tests positive, she is more likely to not have breast cancer.

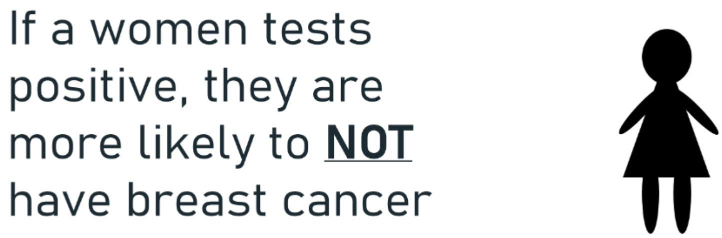

Of course, when I first asked you about the probability that Sally has breast cancer, I did not tell you the base rate.

Regardless, the whole point of this example is to show that we simply cannot tell the probability that someone has breast cancer if we only know the probability of them getting a positive test if they have breast cancer.

Instead, we need more information to determine the probability that someone has cancer given a positive test result. And more specifically, what we need is information about:

1. The base rate, such as how many women have breast cancer in the first place
2. The probability of the test result if they do not have the disease
3. The probability of the test result if they do have the disease

If we do not consider base rates when we should, then we commit the base rate fallacy. So please, avoid the base rate fallacy!

Thank you for watching this video. Please answer the questions below this video and carry on with the study.

References

Berner, E. S., & Graber, M. L. (2008). Overconfidence as a Cause of Diagnostic Error in Medicine. *American Journal of Medicine*, *121*(5 SUPPL.), S2–S23. https://doi.org/10.1016/j.amjmed.2008.01.001

Meyer, A. N. D., Payne, V. L., Meeks, D. W., Rao, R., & Singh, H. (2013). Physicians' diagnostic accuracy, confidence, and resource requests: A vignette study. *JAMA Internal Medicine*, *173*(21), 1952–1961. https://doi.org/10.1001/jamainternmed.2013.10081

Naguib, M., Brull, S. J., Hunter, J. M., Kopman, A. F., Fülesdi, B., Johnson, K. B., & Arkes, H. R. (2019). Anesthesiologists' Overconfidence in Their Perceived Knowledge of Neuromuscular Monitoring and Its Relevance to All Aspects of Medical Practice. *Anesthesia & Analgesia*, *128*(6), 1118–1126. https://doi.org/10.1213/ANE.0000000000003714

Perel, A., Saugel, B., Teboul, J. L., Malbrain, M. L. N. G., Belda, F. J., Fernández-Mondéjar, E., Kirov, M., Wendon, J., Lussmann, R., & Maggiorini, M. (2016). The effects of advanced monitoring on hemodynamic management in critically ill patients: a pre and post questionnaire study. *Journal of Clinical Monitoring and Computing*, *30*(5), 511–518. https://doi.org/10.1007/s10877-015-9811-7

Tetlock, P. (2005). *Expert political judgment: how good is it? How can we know?* Princeton University Press.