

representatives from the major phytoplankton classes in the ocean—diatoms, dinoflagellates, and cyanobacteria—can also produce extracellular superoxide (6, 9, 10). Moreover, field studies have found elevated superoxide concentrations in areas of high phytoplankton abundance (5, 7). Hence, it is now accepted that phytoplankton are the main source of particle-associated superoxide in the upper, photic, oceanic water column (see the figure).

Diaz *et al.* show that extracellular production of superoxide is widespread among taxonomically divergent heterotrophic bacteria from a range of different environments. Some of their bacterial cultures are marine isolates; these bacteria can potentially generate superoxide in marine sediments and in the vast expanses of the deep ocean that do not receive sunlight. Of course, heterotrophic bacteria are not restricted to the deep ocean and may thus also contribute to particle-associated biological superoxide production close to the ocean surface (see the figure).

Superoxide interacts with many chemical elements and compounds. For example, it alters the redox states of iron, copper, and manganese and modulates their chemical reactivity, solubility, bioavailability, and toxicity (8, 9, 13, 14). These metals control the abundance and distribution of marine phytoplankton, which in turn drive the cycling of

major nutrients, such as carbon and nitrogen. Superoxide also oxidizes dissolved manganese to solid manganese oxides, which are efficient trace metal sorbents and powerful oxidants of organic materials (12). When these minerals settle out of the water column, they influence the distribution of trace elements and nutrients. Furthermore, superoxide promotes the degradation of dissolved organic matter, with implications for the marine carbon cycle. Further interactions and biogeochemical roles of superoxide in the ocean are likely.

Given its functions in other systems, superoxide may play a role in the chemical interactions among microorganisms at sea. Superoxide is potentially toxic to organisms and can be used as a first line of defense against viral or bacterial attacks. At low levels, it may also assist communication among marine microbes. So far, the only demonstrated role of superoxide production by phytoplankton is of increased iron availability, shown for a filamentous cyanobacterium (14). However, another study with a diatom found that iron acquisition was unaffected by superoxide production (9).

We are still a long way from a full assessment of superoxide concentrations across oceanic environments and their link to bacterial activity. Given the potential influence of superoxide on trace metal and carbon cycling

in the ocean, these are exciting times to study the dynamics of superoxide in seawater. The analytic capabilities exist, correspondence with other disciplines provides a good stream of ideas and hypotheses, and there are still more questions than answers.

#### References and Notes

1. J. M. Diaz *et al.*, *Science* **340**, 1223 (2013); 10.1126/science.1237331.
2. R. M. Baxter, J. H. Carey, *Nature* **306**, 575 (1983).
3. E. Micinski, L. A. Ball, O. C. Zafriou, *J. Geophys. Res.* **98**, 2299 (1993).
4. S. P. Hansard, A. W. Vermilyea, B. M. Voelker, *Deep Sea Res. J.* **57**, 1111 (2010).
5. A. L. Rose, A. Godrant, M. Furnas, T. D. Waite, *Limnol. Oceanogr.* **55**, 1521 (2010).
6. A. L. Rose, E. A. Webb, T. D. Waite, J. W. Moffett, *Environ. Sci. Technol.* **42**, 2387 (2008).
7. S. A. Rusak, B. M. Peake, L. E. Richard, S. D. Nodder, W. J. Cooper, *Mar. Chem.* **127**, 155 (2011).
8. Y. Shaked, R. Harris, N. Klein-Kedem, *Environ. Sci. Technol.* **44**, 3238 (2010).
9. A. B. Kustka, Y. Shaked, A. J. Milligan, D. W. King, F. M. M. Morel, *Limnol. Oceanogr.* **50**, 1172 (2005).
10. J.-A. Marshall, M. de Salas, T. Oda, G. Hallegraeff, *Mar. Biol.* **147**, 533 (2005).
11. E. Saragosti, D. Tchernov, A. Katsir, Y. Shaked, *PLoS ONE* **5**, e12508 (2010).
12. D. R. Learman, B. M. Voelker, A. I. Vazquez-Rodriguez, C. M. Hansel, *Nat. Geosci.* **4**, 95 (2011).
13. S. P. Hansard, H. D. Easter, B. M. Voelker, *Environ. Sci. Technol.* **45**, 2811 (2011).
14. A. L. Rose, *Front. Microbiol.* **3**, 124 (2012).

**Acknowledgments:** Supported by Israel Science Foundation grant 248/11 (Y.S.).

10.1126/science.1240195

## MATHEMATICS

# Bayes' Theorem in the 21st Century

Bradley Efron

The term “controversial theorem” sounds like an oxymoron, but Bayes' theorem has played this part for two-and-a-half centuries. Twice it has soared to scientific celebrity, twice it has crashed, and it is currently enjoying another boom. The theorem itself is a landmark of logical reasoning and the first serious triumph of statistical inference, yet is still treated with suspicion by most statisticians. There are reasons to believe in the staying power of its current popularity, but also some signs of trouble ahead.

Here is a simple but genuine example of Bayes' rule in action (see sidebar) (1). A physicist couple I know learned, from sonograms, that they were due to be parents of twin boys.

They wondered what the probability was that their twins would be identical rather than fraternal. There are two pieces of relevant evidence. One-third of twins are identical; on the other hand, identical twins are twice as likely to yield twin boy sonograms, because they are always same-sex, whereas the likelihood of fraternal twins being same-sex is 50:50. Putting this together, Bayes' rule correctly concludes that the two pieces balance out, and that the odds of the twins being identical are even. (The twins were fraternal.)

Bayes' theorem is thus an algorithm for combining prior experience (one-third of twins are identicals) with current evidence (the sonogram). Followers of Nate Silver's FiveThirtyEight Web blog got to see the rule in spectacular form during the 2012 U.S. presidential campaign: The algorithm updated prior poll results with new data on

Bayes' theorem plays an increasingly prominent role in statistical applications but remains controversial among statisticians.

a daily basis, correctly predicting the actual vote in all 50 states. “Statisticians beat pundits” was the verdict in the press (2).

Bayes' 1763 paper was an impeccable exercise in probability theory. The trouble and the subsequent busts came from overenthusiastic application of the theorem in the absence of genuine prior information, with Pierre-Simon Laplace as a prime violator. Suppose that in the twins example we lacked the prior knowledge that one-third of twins are identical. Laplace would have assumed a uniform distribution between zero and one for the unknown prior probability of identical twins, yielding 2/3 rather than 1/2 as the answer to the physicists' question. In modern parlance, Laplace would be trying to assign an “uninformative prior” or “objective prior” (2), one having only neutral effects on the output of Bayes' rule (3). Whether or not this

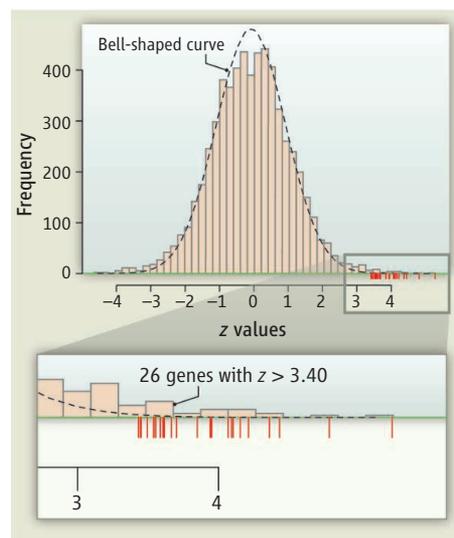
Department of Statistics, Stanford University, Stanford, CA 94305, USA. E-mail: brad@stat.stanford.edu

can be done legitimately has fueled the 250-year controversy.

Frequentism, the dominant statistical paradigm over the past hundred years, rejects the use of uninformative priors, and in fact does away with prior distributions entirely (1). In place of past experience, frequentism considers future behavior. An optimal estimator is one that performs best in hypothetical repetitions of the current experiment. The resulting gain in scientific objectivity has carried the day, though at a price in the coherent integration of evidence from different sources, as in the FiveThirtyEight example.

The Bayesian-frequentist argument, unlike most philosophical disputes, has immediate practical consequences. Consider that after a 7-year trial on human subjects, a research team announces that drug A has proved better than drug B at the 0.05 significance level. Asked why the trial took so long, the team leader replies “That was the first time the results reached the 0.05 level.” Food and Drug Administration (FDA) regulators reject the team’s submission, on the frequentist grounds that interim tests of the data, by taking repeated 0.05 chances, could raise the false alarm rate to (say) 15% from the claimed 5%.

A Bayesian FDA regulator would be more forgiving. Starting from a given prior distribution, the Bayesian posterior probability of drug A’s superiority depends only on its final evaluation, not whether there might have been earlier decisions. This is a corollary of



**True and false discoveries.** Test statistic  $z$  for 6033 genes in a microarray study of prostate cancer. The 28 genes having  $z \geq 3.40$  are likely to be “true discoveries,” that is, genes that are more active in prostate cancer patients than in controls. These results are based on Bayes’ rule, but with “prior” information obtained from the current data, an example of empirical Bayes methodology.

#### SIDEBAR: BAYES’ THEOREM IN ACTION

If  $P(A)$  is the probability of A and  $P(B)$  is the probability of B, then the conditional probability of A given B is  $P(A|B)$  and the conditional probability of B given A is  $P(B|A)$ . Bayes’ theorem says that

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

In the twins example, A is “twins identical” and B is “sonogram shows twin boys.” The doctor’s prior says  $P(A) = 1/3$ ; genetics implies  $P(B|A) = 1/2$  and  $P(B|\text{not } A) = 1/4$ , so  $P(B) = (1/2)(1/3) + (1/4)(2/3) = 1/3$ . Bayes’ theorem then gives

$$P(A|B) = (1/2)(1/3)/(1/3) = 1/2$$

The two pieces of evidence thus balance out, and the likelihood of the boys being fraternal is equal to that of the boys being identical.

Bayes’ theorem, convenient but potentially dangerous in practice, especially when using prior distributions not firmly grounded in past experience.

I recently completed my term as editor of an applied statistics journal. Maybe a quarter of the papers used Bayes’ theorem. Almost all of these were based on uninformative priors, reflecting the fact that most cutting-edge science does not enjoy FiveThirtyEight-level background information. Are we in for another Bayesian bust?

Arguing against this is a change in our statistical environment. Modern scientific equipment pumps out results in fire hose quantities, producing enormous data sets bearing on complicated webs of interrelated questions. In this new scientific era, the ability of Bayesian statistics to connect disparate inferences counts heavily in its favor.

An example will help here. In a microarray prostate cancer study (4), 102 men—52 patients and 50 healthy controls—each had their genetic activity measured for 6033 genes. The investigators were hoping to find genes expressed differently in the patients than in the controls. To this end, they calculated a test statistic  $z$  for each gene, with a standard normal (“bell-shaped”) distribution in the null case of no patient/control difference, but with bigger values for genes expressed more intensely in patients.

The histogram of the 6033  $z$  values (see the figure) does not look much different than the bell-shaped curve that would apply if all genes were null. However, there is a suggestion of interesting non-null genes in the heavy right tail of the distribution. We have to be careful, though. With 6033 genes to consider at once, a few of the  $z$ ’s are bound to look big even under the null hypothesis, an example

of selection bias or regression to the mean. These would be “false discoveries.”

False discovery rates (FDRs) (5) are a recent development that takes multiple testing into account (6). Here, it implies that the 28 genes with  $z$  values above 3.40 (red dashes in the figure) are indeed interesting, with the expected proportion of false discoveries among them being less than 10%. This is a frequentist 10%: how many mistakes we would average using the algorithm in future studies. We expect only 2.8 of the  $z$  values exceeding 3.40 to be null, that is, only 10% of the actual number observed. Larger choices of the cutoff would yield smaller FDRs.

This brings us back to Bayes. Another interpretation of the FDR algorithm is that the Bayesian probability of nullness given a  $z$  value exceeding 3.40 is 10%. What prior evidence are we using? None, as it turns out! With 6033 parallel situations at hand, we can effectively estimate the relevant prior from the data itself. “Empirical Bayes” is the name for this sort of statistical jujitsu, suggesting a fusion of frequentist and Bayesian reasoning (7). Empirical Bayes is an exciting new statistical idea, well-suited to modern scientific technology, saying that experiments involving large numbers of parallel situations carry within them their own prior distribution. The idea was coined in the 1950s (8), but real developmental interest awaited the vast data sets of the 21st century.

I wish I could report that this resolves the 250-year controversy and that it is now safe to always employ Bayes’ theorem. Sorry. My own practice is to use Bayesian analysis in the presence of genuine prior information; to use empirical Bayes methods in the parallel cases situation; and otherwise to be cautious when invoking uninformative priors. In the last case, Bayesian calculations cannot be uncritically accepted and should be checked by other methods, which usually means frequentistically.

#### References and Notes

1. B. Efron, *Bull. Am. Math. Soc.* **50**, 129 (2013).
2. S. Wang, B. Campbell, *Science* **339**, 758 (2013).
3. J. Berger, *Bayesian Anal.* **1**, 385 (2006).
4. D. Singh et al., *Cancer Cell* **1**, 203 (2002).
5. Y. Benjamini, Y. Hochberg, *J. R. Stat. Soc. B* **57**, 289 (1995).
6. See chapter 4 of (7) for a careful exposition of false discovery rate theory.
7. B. Efron, *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction* (Institute of Mathematical Statistics Monographs, Cambridge Univ. Press, Cambridge, UK, 2010).
8. H. Robbins, in *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955, Vol. I* (Univ. of California Press, Berkeley/Los Angeles, 1956), pp. 157–163.