

EVERYTHING

[Nick Bostrom (2008): Response to the *2008 EDGE Question*: “What Have You Changed Your Mind About?”]

www.nickbostrom.com

For me, belief is not an all-or-nothing thing—believe or disbelieve, accept or reject. Instead, I have degrees of belief, a subjective probability distribution over different possible ways the world could be. This means I am constantly changing my mind about all sorts of things, as I reflect or gain more evidence. While I don’t always think explicitly in terms of probabilities, I often do so when I give careful consideration to some matter. And when I reflect on my own cognitive processes, I must acknowledge the graduated nature of my beliefs.

The commonest way in which I change my mind is by concentrating my credence function on a narrower set of possibilities than before. This occurs every time I learn a new piece of information. Since I started my life knowing virtually nothing, I have changed my mind about virtually everything. For example, not knowing a friend’s birthday, I assign a $1/365$ chance (approximately) of it being August 11. After she tells me that August 11 is her birthday, I assign that date a probability of close to 100 percent. (Never exactly 100 percent, for there is always a non-zero probability of miscommunication, deception, or other error.)

It can also happen that I change my mind by smearing out my credence function over a wider set of possibilities. I might forget the exact date of my friend’s birthday but remember that it is sometime in the summer. The forgetting changes my credence function from being almost entirely concentrated on August 11 to being spread out more or less evenly over the summer months. After this change of mind, I might assign a 1-percent probability to my friend’s birthday being August 11.

My credence function can become more smeared out not only by forgetting but also by learning—learning that what I previously took to be strong evidence for some hypothesis is in fact weak or misleading evidence. (This type of belief change can often be mathematically modeled as a narrowing rather than a broadening of credence function, but the technicalities of this are not relevant here.) For example, over the years I have become moderately more uncertain about the benefits of

medicine, nutritional supplements, and much conventional health wisdom. This belief change has come about as a result of several factors. One of the factors is that I have read some papers that cast doubt on the reliability of the standard methodological protocols used in medical studies and their reporting. Another factor is my own experience of following up on MEDLINE some of the exciting medical findings reported in the media—almost always, the search of the source literature reveals a much more complicated picture, with many studies showing a positive effect, many showing a negative effect, and many showing no effect. A third factor is the arguments of a health economist friend of mine, who takes a dim view of the marginal benefits of medical care.

Typically, my beliefs about big issues change in small steps. Ideally, these steps should approximate a random walk, like the stock market. It should be impossible for me to predict how my beliefs on some topic will change in the future. If I believed that a year hence I would assign a higher probability to some hypothesis than I do today—why, in that case, I could raise the probability right away. Given knowledge of what I will believe in the future, I would defer to the beliefs of my future self, provided I think my future self will be better informed than I am now and at least as rational.

I have no crystal ball to show me what my future self will believe. But I do have access to many other selves, who are better informed than I am on many topics. I can defer to experts. Provided they are unbiased and are giving me their honest opinion, I should perhaps always defer to people who have more information than I do—or to some weighted average of expert opinion if there is no consensus. Of course, the proviso is a big one: Often I have reason to disbelieve that other people are unbiased or that they are giving me their honest opinion. However, it is also possible that I am biased and self-deceiving. An important unresolved question is how much epistemic weight a wannabe Bayesian thinker should give to the opinions of others. I'm looking forward to changing my mind on that issue, hopefully by my credence function becoming concentrated on the correct answer.