When NOT to Update Your Beliefs

Synopsis
I argue that in cases with low prior probabilities and unreliable evidence (e.g. personal anecdotes), it is rational not to update one's posterior probabilities at all in response to additional low quality evidence (e.g. an additional anecdote). I present my basic case with reference to Bayes' Theorem, and then consider some rebuttals. I reject that rebuttal that updates should be small but non-zero on the grounds that such small updates are within the bounds of error of one's probabilities. I reject the rebuttal that many anecdotes provide stronger cumulative evidence on the basis that anecdotes are not independent events. I conclude with a discussion about the differences between updating in abstract theory, and updating in the real world.

The Basic Argument

Let $P(\text{Aliens})$ be 'the probability that aliens exist and are currently visiting the Earth'
Let $P(\text{Report})$ be 'the probability that a given report of seeing or being abducted by aliens is made'

From Bayes' Theorem we have:

$$\text{P(Aliens|Report)} = \frac{\text{P(Report|Aliens)} \times \text{P(Aliens)}}{\text{P(Report)}}$$

Dividing both sides by the prior to focus attention only on the 'update' quotient we have:

$$\frac{\text{P(Aliens|Report)}}{\text{P(Aliens)}} = \frac{\text{P(Report|Aliens)}}{\text{P(Report)}}$$

Expanding using marginal probabilities we have:

$$\frac{\text{P(Aliens|Report)}}{\text{P(Aliens)}} = \frac{\text{P(Report|Aliens)}}{\text{P(Report|Aliens)} \times \text{P(Aliens)} + \text{P(Report|}{\sim}\text{Aliens)} \times \text{P(}{\sim}\text{Aliens)}}$$

Let us suppose that if aliens do exist, then reports of alien abduction are certain (note this assumption is not crucial and is only for simplicity), hence:

$$\text{P(Report|Aliens)} = 1$$

$$\frac{\text{P(Aliens|Report)}}{\text{P(Aliens)}} = \frac{1}{\text{P(Aliens)} + \text{P(Report|}{\sim}\text{Aliens)} \times \text{P(}{\sim}\text{Aliens)}}$$

Suppose P(Aliens) is very low (i.e. we have a low prior), such that:

$$\frac{\text{P(Aliens|Report)}}{\text{P(Aliens)}} \approx \frac{1}{\text{P(Report|}{\sim}\text{Aliens)}}$$

The question of how much to update one's belief in aliens thus comes down to what value is given to $\text{P(Report|}{\sim}\text{Aliens)}$. Suppose that, for a variety of reasons (such as the many biases of human memory and perception, and the great number of historical examples of these sorts of incredible stories), I think that $\text{P(Report|}{\sim}\text{Aliens)}$ is high, close to 1. That is, I expect reports of aliens to be very likely even if aliens don't exist, because this form of evidence is highly unreliable. In this case, we have:

$$\frac{\mathrm{P(Aliens|Report)}}{\mathrm{P(Aliens)}} \approx 1$$

Put a different way:

$$\mathrm{P(Aliens|Report)} \approx \mathrm{P(Aliens)}$$

This effectively means that no update is made: posterior probability is approximately equal to the prior probability.

'You are Ignoring the Delta' Rebuttal

The natural rebuttal to this argument is to say "approximately equal doesn't mean exactly equal. You should still update, even if not very much". I disagree, at least when it comes to actual real-world beliefs and not the ideal platonic realm of theory. This is because, in the real world, probabilities are often not known with precision, but can generally only be given with considerable error bars.

Let me give an example. In the real world, I don't know what $\mathrm{P(Aliens)}$ is. I am uncertain, so I give it a probability distribution with fairly large variance. Let me include some illustrative but arbitrary numbers:

$$\mathrm{P(Aliens)} = 0.010 \pm 0.005$$
$$\mathrm{P(Report|{\sim}Aliens)} = 0.96$$

Hence:

$$\frac{\mathrm{P(Aliens|Report)}}{\mathrm{P(Aliens)}} = \frac{1}{0.01 + 0.96 \times 0.99}$$
$$\frac{\mathrm{P(Aliens|Report)}}{\mathrm{P(Aliens)}} = 1.041$$

$$\mathrm{P(Aliens|Report)} = 0.01041 \pm 0.004$$

What we observe is that the change in my belief as a result of updating on the report is very small compared to the existing uncertainty on my belief. Indeed, generally it is not meaningful to report estimates with more significant figures than are implied by the error bars. As the error bars in this example are so much larger than the updated estimate, the new confidence interval is simply identical to the original confidence interval:

Before:

$$\mathrm{P(Aliens)} = \{0.005{,}0.015\}$$

After:

$$\mathrm{P(Aliens)} = \{0.005{,}0.015\}$$

There are many ways of expressing this idea: the information provided by the report is so small that it is swamped by existing uncertainties, the amount of information provided by the report is so small that it not worth the effort to update my beliefs on it, the amount of information provided by the report is so small that it will never realistically effect any other actions or beliefs that I care about.

'Many Anecdotes are Different' Rebuttal

Another potential rebuttal to my argument is to say that, while this might work for a single case, many reports of aliens would have a much larger effect on our beliefs. In other words, it may be the case that $\delta \approx 0$, but that doesn't mean that $1000\delta \approx 0$. In my view, however, this rebuttal too fails, as it erroneously assumes that the probabilities of multiple reports are independent.

Consider the probability that a report is made given that aliens don't exist:

$$\mathrm{P(Report|{\sim}Aliens)} = p$$

Now consider that we know of one case, $\mathrm{Report_1}$, when a known false report was made, and a second case, $\mathrm{Report_2}$. It seems very reasonable to assert that:

$$\mathrm{P(Report_2|Report_1, {\sim}Aliens)} > p$$

In other words, if we already have one case of a false report, then the probability of another false report increases. Now consider the following situation of many such false reports for large n:

$$\mathrm{P(Report_n|Report_1, Report_2, \dots, Report_n, {\sim}Aliens)} \gg p$$

If we consider again our update equation:

$$\frac{\mathrm{P(Aliens|Report)}}{\mathrm{P(Aliens)}} = \frac{1}{\mathrm{P(Aliens)} + \mathrm{P(Report|{\sim}Aliens)} \times \mathrm{P({\sim}Aliens)}}$$

A large value of $\mathrm{P(Report|{\sim}Aliens)}$ will amplify the effect of a large $\mathrm{P({\sim}Aliens)}$, thus producing an update proportion of very close to 1 (i.e. a very small update).

Thus we see that, given the presence of many past instances of false reports, it is rational to respond to each additional report with smaller and smaller updates. This effect would be further magnified if we considered additional complexities, such as the fact that we have good reason to suppose that false reports are the product of common psychological and cultural processes (and hence are even less likely to be independent), but I will not go through a detailed analysis of that here.

Uncertainty and Pragmatics

I agree that in an ideal platonic realm of infinite processing power, fully known probabilities, and excellent knowledge of the conditionality of various probabilities (etc), one should simply plug all the numbers into Bayes' Theorem and update exactly that much. What I am arguing is that, in the messy real world, when probabilities are unknown, when the exact conditional relationships between events are not known, and importantly, when processing power and time are limited, the rational thing to do is to adopt useful heuristics about when and how much to update one's beliefs.

I have argued in this piece that one of these heuristics should be not to update beliefs when (1) priors are low and (2) evidence is of poor quality. I believe that any attempt to update beliefs in these circumstances will (1) not be worth the effort, (2) produce updates that are so small they are meaningless compared to uncertainties in the probability estimates themselves, (3) probably lead to erroneously large updates on the basis of failing to consider all relevant factors of conditional probabilities and reliability of the evidence (which is almost impossible to do in practise for highly equivocal evidence).