Useful steps when thinking about any difficult or disconcerting issue are to state the components of the issue as clearly as possible, and combine them in a way that is logically and mathematically justified.

Such an issue is the recent claim that an 1816 scriptural-style history of the War of 1812 entitled *The Late War between the United States and Great Britain* (*LW*) had an influence on the Book of Mormon. Since "influence" can mean many things, the claim in and of itself isn't necessarily disconcerting to believers in the Book of Mormon. After all, whatever else the translation process involved (divine inspiration, angels, plates, interpreters, stones, hats, scribes), it involved Joseph Smith's mind. Joseph Smith's mind and language were immersed in the same culture that produced *LW*. It would have been very difficult for Joseph Smith, even as translator, to express concepts without reference to his culture. Furthermore, the purpose of scripture is to impart timeless lessons. If some of Joseph Smith's cultural milieu made its way into the translation of the Nephite record as Joseph recognized obvious lessons from antiquity for himself and his times, I see nothing disconcerting.

The recent claim of influence of *LW* on the Book of Mormon, however, was deliberately intended to be troubling to believers because the paper presenting the claim was provocatively entitled 'How the Book of Mormon Destroyed Mormonism.' Almost immediately after presentation of the paper, posters on several comment boards joined in, describing the claim as the 'Smoking Gun' or the 'Silver Bullet' that would bring down Mormonism.

Support for the claim was based on a numerical measure of similarity devised by the authors of the paper. Although the measure was an *ad hoc* index not based on a formal model of English composition, it was based on the sensible idea that shared rare phrases should be more indicative of similarity than shared common phrases. The measure ignored other important things, however, like relative frequencies of matching phrases. When the measure of similarity was calculated between the Book of Mormon and each of 5000 books randomly selected from a large corpus written between 1500 and 1830, *LW* emerged as most similar to the Book of Mormon after removal of biblical phrases.

LW does have interesting parallels to the Book of Mormon, but Ben McGuire, in an excellent discussion in the Interpreter Foundation's journal [https://interpreterfoundation.org/the-late-war-against-the-book-of-mormon/], pointed out troubling issues with the similarity measure, problems with the corpus selected for comparison, problems with the process of cleaning the texts and extracting phrases, and interpretational problems based on the simple fact that *LW* was identified as a result of a massive search. Some critics of the church acknowledged Ben's thoughtful analysis, but dismissed it, saying that statistical results can always be debated and nothing ever proved via statistics. One critic said that he could plainly see the influence of *LW* on the Book of Mormon upon a close reading of the texts; he wished he had the skill to devise an appropriate statistical formula to calculate a probability that would express what he thought he saw. This short paper attempts to give this critic as well as believers a formula to do the job.

To formalize things, I start by defining some components:

1. Let H be the (hypothetical) proposition that the production of the Book of Mormon was influenced by *LW*. We immediately notice the need to be specific about what is meant by "influenced." I suggest that we take "influenced" to mean "was deliberately and physically used by Joseph Smith in his composition of the Book of Mormon." We do not have to go this far in the critical direction in formulating this proposition, but we do have to be very specific and we have to insist that the definition not change in the course of discussion. So, let's start with

this. Also, let H* be the opposite proposition of H.

2. Let E be the observation that, after an extensive search for matching phrases in 5000 randomly chosen books, several hundred non-biblical matches between *LW* and the Book of Mormon are discovered. Using a weighting system for matching phrases in which weights decrease severely (actually, hyperbolically) as frequencies of occurrence of phrases in the corpus increase, *LW* turns out to have the highest similarity score of all 5000 books in the corpus.
3. A person's belief in H before encountering E can be expressed as a probability P(H), called the *prior probability of* H. If someone has no opinion about H, but wants to use probability theory as a mind game for thinking about potential results, it is convenient to set P(H)=0.5. The person's disbelief in H before encountering E is denoted by P(H*)=1-P(H), which is also 0.5 when P(H)=0.5.
4. P(H|E) is the modified (or *posterior*) probability of proposition H given that E has occurred. Our objective is to evaluate P(H|E) if possible – or at least to discuss it sensibly. P(H|E) must not be confused with the reverse probability P(E|H), the probability of evidence E occurring if H is true. In general, P(E|H) is straightforward to think about while P(H|E) is hard *if not impossible* to evaluate naively.

The Bayes rule, however, is a relatively straightforward probabilistic formula that indicates, under broad conditions, how to properly calculate P(H|E) from P(E|H) , P(E|H*), and P(H). The Bayes rule summarizes the thinking that has to go on when reasoning from evidence (E) to a proposition (H).

Results using the Bayes rule can be counterintuitive. Here is a classic example. The proposition (H) is that a random 40-year-old woman has breast cancer. The incidence of breast cancer in the U.S. among 40-year-old women is 0.0144. Hence the *prior probability* P(H) is 0.0144. A mammogram is a good diagnostic test. It has probability P(E|H)=0.84 of being positive when the subject has cancer, and probability P(E|H*)=0.10 of falsely being positive when the subject does not have cancer. A random woman decides to have a mammogram. The test is positive. The patient is terrified, but the critical quantity is P(H|E), the *posterior probability* that the woman has breast cancer given that the mammogram is positive. It would seem that with such a precise diagnostic test, P(H|E) would be very high. However, it turns out that P(H|E) is only 0.11 = 11%. This example graphically illustrates how bad our intuition can be when reasoning about a proposition given evidence. The value 0.11 was calculated as follows by the Bayes rule:

$$P(H|E) = \frac{P(E|H)P(H)}{P(E|H)P(H) + P(E|H^*)P(H^*)} = \frac{(0.84)(0.0144)}{(0.84)(0.0144) + (0.10)(0.9856)} = 0.11.$$

It is easy to get lost in the reasoning of any argument, even an argument that we ourselves are formulating. As a statistician, and therefore being lazy but interested in the truth, I try to offload the work of formulating and evaluating a complex argument, as much as possible, onto the laws of mathematics. The Bayes rule is one of the most versatile and useful of such laws.

To apply the Bayes Rule to the *LW* – Book of Mormon situation, I propose that we set the priors P(H) and P(H*) to 0.5 to express the idea that we are willing to suspend our innate feelings about the Book of Mormon one way or the other, at least for the sake of argument. Since these proposed values are equal (P(H) = P(H*) = 0.5), the Bayes rule simplifies to

$$P(H|E) = \frac{P(E|H)}{P(E|H) + P(E|H^*)}.$$

In my opinion, P(E|H) (= the probability of the evidence assuming that *LW* is a direct source for the Book of Mormon) is at most 0.8. Even if *LW* is a source for the Book of Mormon, a book other than *LW* may well be identified as more similar to the Book of Mormon for several reasons: 1. since the corpus was chosen at random, the weights and therefore similarity results will differ by corpus, 2. the population from which the corpus was selected might tend to not generate optimal weights for choosing the Book of Mormon, 3. the ever-present possibility of unintended parallelisms (especially in the massive search model) between two completely unrelated books might affect identification of *LW*, 4. a similarity measure that ignores frequencies of use of common phrases or does not adjust properly for lengths of the texts might tend to not identify *LW*, and 5. books in the corpus may not have been text-cleaned properly leading to over- or under-valuation of the true similarity of *LW* to the Book of Mormon and other books. Very little information is available on the effects of these 5 issues on P(E|H), so it is premature to assume that P(E|H) is too large.

In my opinion, P(E|H*) (= the probability of the evidence assuming that *LW* is *not* a direct source for the Book of Mormon) is at least 0.5. Some reasons why *LW* might have a high similarity measure even though it had no influence on the Book of Mormon are: 1. the massive search model might identify unintended parallels in books not connected with the Book of Mormon, 2. unrelated books that use deliberate scriptural language will share many phrases with the Book of Mormon, 3. common phrases might appear for unrelated books that arise in similar cultural settings, 4. common segments such as the copyright statement might induce apparent similarity between two unrelated books, 5. properties of the *ad hoc* similarity index might accidentally favor certain books in the corpus, and 6. in any corpus there will always be a 'most similar' book to the Book of Mormon, even if none of the books in the corpus had anything to do with the Book of Mormon. *LW* might simply have been that book.

Given these considerations, Bayes rule calculations give:

$$P(H|E) = \frac{P(E|H)}{P(E|H) + P(E|H^*)} = \frac{0.8}{0.8 + 0.5} = .62.$$

Hence P(H|E) increases at most to around 0.6 from 0.5. That is, the occurrence of E (= *LW* has the highest similarity score of all 5000 books in the corpus) increases the likelihood of H (= a direct influence of *LW* on the Book of Mormon) from a 50-50 situation to a 60-40 situation. No slam dunk has occurred. You can argue with my reasoning about the values of P(E|H), P(E|H*) and P(H), but now the arguments are focused on real issues that can be debated and investigated carefully.

I recommend the use of the Bayesian framework for discussing the significance of authorship questions using corpus-based similarity statistics. The framework is also useful for determining what future work is needed to clarify the issues and allowing more exact evaluations of the relevant quantities.