The Theory That Would Not Die

By Sharon Bertsch McGrayne

An Enhanced Reading Guide by **Farnam Street**

Table of Contents

Introduction

For four weeks in January of 2018, a large group of bright and attentive Farnam Street members read the book *The Theory That Would Not Die* by Sharon Bertsch McGrayne together and discussed a set of questions about each week's reading selection.

This document is a collation of those questions, and a sampling of the most insightful comments each week, intended to complement to your own reading of *The Theory That Would Not Die*.

Week 1

Reading covered: Chapters 1-3 of The Theory That Would Not Die.

Question One
Does the logic underlying Bayesian thinking & probability estimating feel logical, intuitive - is it the way you already reason - or is it slightly challenging? Can you think of any uses for this style of judgment yet?

> *It does seem intuitive and that is both its strength and its risk/danger. I speak from the perspective of a scientist - we are trained and actively discouraged from taking our intuition into account and to leave our biases aside. (I am not saying we always, or even mostly succeed - but most of us work hard at it.) So saying, make an educated guess and then gather data sounds dangerous. On the other hand, with equal priors perhaps my main concern is aleviated.*
>
> ***
>
> *Intellectually, it seems very familiar (...obvious?), habitually, the challenge seems to be to weigh Prior and the new evidence appropriately. I find that new data can lead to salience bias, which can lead to disproportionately ignoring why you established the Prior. I seem to remember that's something discussed in Superforecasters, and it's a personal challenge. Philosophically, I find the problem of induction a real and continuous thorn for any forecasting, including Bayes (assuming that, ultimately, even finding prior causes is intended to ultimately help build a better model to forecast).*
>
> ***
>
> *I come to think about social relationships as an area where I apply the Bayesian style of judgement.*

When interested in another person, either for romance, for friendship or for business, I start with a 50/50 chance that the person is interested, try different attempts to check if the estimate is probable, and update my belief along the way (sometimes this takes just a couple of minutes, sometimes it takes a lifetime).

***

I agree with @bjorke about the lack of detail, especially with respect to the Dreyfus trial. The reasoning behind the use of Bayesian statistics to prove forgery sounds novel and counterintuitive, which sounds like mental map territory. As a former mortar-man, I would have enjoyed more detail on the use of statistics in artillery fire: why are Bayesian tables more efficient than bracketing fire and trigonometry?

All that being equal, the logic behind Bayes' (-Laplace-Price's) theorem seems intuitive. My brother-in-law and I have been locked in a discussion about the shortcomings of the predominant methods of using power to quantify physiological responses in cyclists. (A mutual hobby.) He tends to take exception with the methods because they can't account for initial conditions; my view is that initial conditions can't really be accounted for, and that the level of uncertainty in these methods are _good enough_ for the predicted response. Sound familiar?

***

I live in the work of statistics but I'm industry (as opposed to academic). From that I feel that while reading about Bayesian thinking seems intuitive it is not. All the research points to humans NOT being natural statisticians. We tend to make conclusions quickly and not look let go of them. We need to be trained to be more open minded and logical. Especially when faced with little or limited data holding beliefs loosely would be valuable but human nature is the opposite to make a firm conclusion and not change.

***

What struck me about the "frequentist" vs. "Bayesian" argument is how hard it is to say that the two shall never meet when the primary developer of Bayes' theory was also the inventor of the Central Limit Theorem. I view the approaches as different tools in a toolbox; from that standpoint, it's difficult to see these disputes as anything more significant than groups of people angry at other groups of people for pointing out that their circular saw won't hammer nails.

***

Without ever associating Bayes' or LaPlace's name, I have intuitively applied the thinking to decisions-making and problem-solving. Often, the pitfalls of limiting priors and evaluating new evidence were not so intuitive and a fuller understanding *might have led to more desirable outcomes. It strikes me also that the examples (e.g. Enigma, Bell Telephone) are evaluated historically and in many cases there was scant familiarity with Bayes theories to say they were used knowingly. Without the benefits of computer processing, it doesn't surprise me that Bayes was rejected by decision makers for so many decades. It makes me wonder - what are we missing today?*

***

I think I have been using Bayesian thinking in the course of my work and reading this book has come as a pleasant surprise. I did not know the work methodology has a name for it. My team works with R&D and strategy teams of companies, helping them with competitive strategy, future product/technology roadmaps, and marketing. We often need to consider different scenarios of what led our clients to where they are in the current marketplace and what events/strategies would help them in future. We use different analytical frameworks and information collected through desk research to create an initial set of hypothesis (priors) and then we accept/discard these hypothesis using real-world testing and feedback. Since we largely work with identifying new application areas for existing products and/or exploring under-the-radar technologies, we often need to consider various assumptions as the starting point, which we then seek to validate through further research and feedback.

Question Two
McGrayne gives us a flavor of the early history of Bayesian probability in the first few chapters -- starting with Bayes himself, the great genius Pierre Simon Laplace, and into the 20th Century and the development of modern statistics.

She hints at, but unfortunately never _fully_ describes a debate between "frequentists" and "Bayesians". To simplify (perhaps overly so), the frequentists feel the Bayesians carried a harmful hint of subjectivity: Their "prior probabilities" are randomly selected (especially if we assign equi-probable weights to clearly not equiprobable events, as with the paternity example), or selected on a hunch. The only reasonable conclusions which can be made, said the frequentists, come from examining definite data for a distribution of known outcomes and using statistical methods based on that countable data. Both, of course, have successes to their name - and both are in some ways compatible. (Again - I do wish McGrayne explained this more.)

What do you take of this debate so far? Do you understand it? Why do you think Bayesian approaches to logic might have been slightly cast aside at first?

I found this part very interesting, and I have to confess discovered that I was raised (as a trained scientist, educated in the UK) entirely in the frequentist camp without even knowing there was another. I am a biologist, not a mathematician, and statistics *was a tool we learned that would allow us to examine whether a data set we had generated was meaningful or not. We were actively discouraged from forming an opinion before the statistical analysis, and taught to couch any language in papers or presentations in terms of p values etc. A such Bayes seems a little heretical, but at the same time so useful for real life where you are not going to have controlled experiments and repeat data sets. (Do economists like Bayes?)*

***

It appears to me that the battle of ideas (frequentists vs. Bayesians) stem from deep commitment to objectivity and precision hence anchored in data; and a measure of belief supplemented with other things we can learn from - much broader range between absolute certainty and absolute uncertainty. to be comfortable with the latter - i suppose the formula is: reasonable person+intuition= initial belief you can use and improve with follow on objective information. Early opposers to the Bayesians must have given too much credit to human capacity (or maybe religious context forced them?) to be able to know the causes with certainty.

The phrase I liked in the book so far: 'a person's subjective degree of belief could be measured by the amount he was willing to bet' by Emile Borel.

***

Two things about this argument stand out to me:

1. Every "update" in Bayesian statistics is based on an observation must be grounded in an objective observation of reality. Without going into the philosophical weeds about the nature of reality, if we agree that reality is either objective or a well-distributed delusion, I don't see the conflict between "frequentists" and "Bayesians" as being much more than one more internecine academic argument. They're simply different methods for different approaches.

2. The author portrays it as two camps in a life-or-death struggle, but I have found that life is very rarely that cut-and-dry. I am curious how large this conflict actually was: truly a life-or-death struggle between a bitterly-divided academic community; or a small, pitched battle in one part of statistics? If the latter, does that fact water down the thesis of the book?

As someone who has practically used the Laplace transform a good many times, I'm more than willing to give the guy a shake if one of his methods plays a little fast and loose with statistics. That said, I feel the most important thing in statistics is transparency in reporting error in any statistical study.

***

As with Matt KG and Jonas Blom, I had the same experience in biology grad school as far as only one approach to statistics being taught (Fisherian) and Bayes not even mentioned; I didn't discover it until almost 2 decades later and realized that most of my understanding about statistics from graduate school was wrong, i.e., that a p value tells me absolutely nothing about what I was trying to discover: the validity of my research hypothesis. Given these 4 questions:

1. Given these data, is my research hypothesis (H1) true?
2. Given these data, what is the probability that my research hypothesis (H1) is true?
3. Given these data, what is the probability that the null hypothesis (Ho) is true?
4. Given that the null hypothesis (Ho) is true, what is the probability of these (or more extreme) data?

The only thing that a p value indicates from frequentist statistical tests is an answer to the 4th question, and it says absolutely nothing about the first question, or the answer to what most scientific investigators are trying to find out. The beauty of Bayesian approaches is that they do provide answers to the first question.

One of the points made on page 55 which the author failed to clarify (and which is also misleading in many statistical textbooks) is the following:

"Technically, p-values let laboratory workers state that their experimental outcome offered statistically significant evidence against a hypothesis if the outcome (or a more extreme outcome) had only a small probability (under the hypothesis) of having occurred by chance alone."

NOTE: The clarification needed is that the p-value offered laboratory workers statistically significant evidence against the null hypothesis, which is rarely, if ever, true; therefore, as Carver pointed out in his excellent 1978 review article, statistical significance means little or nothing (Carver, R. P. 1978. The case against statistical significance testing. Harv. Ed. Rev. 48: 378-399)

One of the threads in the book I enjoyed most in the book was the description of Pearson's & Fisher's "kind of feuding and professional bullying generally seen only on middle school playgrounds" (also something I witnessed on a regular basis in the halls of ivy league science departments in both undergraduate & graduate school). It reminded me of Wallace Sayre's famous quip about the "politics of academics being especially bitter because the stakes are so low..

***

fs

The biggest retort from the frequentists is that Bayes is subjective, aka there are assumptions. The problem with this retort is that frequency statistics has as many if not more assumptions, they are just buried in the background. Assumptions of normality of data and errors (which are frequently wrong), assumptions of distributions (that don't match the real world), assumptions of homoscedasticity. And this leads to a false sense of certainty. Many non statisticians using these methods without fully understanding the assumptions of the real interpretation of the outcomes. I go through this at my work trying to explain that a p value of .03 doesn't mean a 3% chance of error. It means if you calculated the mean of many random samples you would be likely to get this mean 3% assuming your null hypothesis is true.

Week 2

Reading covered: Chapters 4-8 of The Theory That Would Not Die.

Question One

We get a hint of where Bayesian and frequentist thinking differ in McGrayne's description of how Bayesian thinking affected practical testing: "In sequential analysis, once several tests or observations strongly cleared or condemned a case of, say, field rations or machine-gun ammunition, the tester could move on to the next box. This almost halved the number of tests required..."

Why is the Bayesian approach is so much faster - more importantly, where else could we apply this type of "sequential" thinking in preference to sample-size dependent thinking?

> *I guess the obvious answer to why it's faster is "it requires fewer observations.", but the reasoning behind _that_ is more complicated. I suppose from a philosophical standpoint, the Bayesian approach views probability as a subject degree of belief about a hypothesis, whereas the Frequentist approach defines probability in terms of a relative frequency in a large number of trials. The latter carries a higher burden of proof.*
>
> *I believe the sequential approach is used in some clinical trial designs. It can also be used for A/B testing to compare effectiveness of two websites or ad campaigns.*
>
> *Note: Abraham Wald, who is referenced in the book, wrote a book on the subject called "Sequential Analysis".*

***

> *When I took undergraduate statistics back in the dark ages, Bayesian approaches weren't taken seriously in part because non-trivial problems were computationally intractable. Timeshare systems, minicomputers, and early PC's had very limited capacity. If you wanted to build models you had to write a fair amount of code. Tukey's work on the 1960 election required a small team and access to one of the most powerful computers then available. Not many of us were so blessed, until perhaps the 1990's.*

***

> *Bayesian folks didn't excel in explaining what their approach was best suited to, i.e. drawing conclusions without the need of well controlled experiments. Similarly, frequentists retreated to a philosophically defensible position that when rigor is*

needed, data cannot substitute anything else. In addition, the academic environment encourages religious-like behaviors about intellectual purity and the custody of truth. I think you got most of the ingredients here. It seems to me that it is not by chance that the Bayesian approach grew sort of outside the academic word.

Question Two

Why do you think leads Bayesian devotees to such a messianic devotion to it - as evidenced by people like Jimmie Savage? What is inside the approach which makes it so seductive - and so potentially divisive - compared to alternative approaches?

I think the "messianic devotion" (at least for people trained in statistics during the 20th century) comes from having the "scales fall from my eyes" when those who have been taught only one way to do statistics (Fisherian approach with null hypothesis statistical testing) are exposed to the Bayesian approach and realize that most of what they were taught about statistics fostered a misunderstanding of what the p value from a statistical test was telling them. The fantasies about Fisherian statistical significance testing (or obtaining a p value of 0.05 or less) fall into 4 categories:

1. The illusion of attaining improbability, or thinking that denying a correct, probabilistic initial premise will result in a sensible conclusion.

2. The odds-against-chance fantasy, thinking that the p value is the probability that the results were caused by chance or that 1-p represents the probability that the results were not caused by chance.
3. The fantasy of research hypothesis validity, that obtaining a p value of 0.05 or less says something about the research hypothesis instead of something about the rareness of the data given that the null hypothesis is true.
4. The fantasy of replicability, that obtaining a p value of 0.05 means that we can be 95% confident that the results are "reliable" or that the probability is 0.95 that the results will replicate.

Neyman & Pearson in their 1933 seminal theoretical work categorized the rejection of the null hypothesis as a function of 5 factors, the last of which was the number of observations; it is the dependency on this last factor that is the ultimate weak link, as Nunnally pointed out in his 1960 paper: "....if the null hypothesis is not rejected, it is usually because the N is too small. If enough data are gathered, the hypothesis will generally be rejected. If rejection of the null hypothesis were the real intention in psychological experiments, there usually would be no need to gather data. (p. 643)".

When one is then exposed to Bayes theorem and realizes that the p value reached by classical Fisherian methods is not a summary of the data or reveal anything about how strong or dependable the particular result is; investigators and readers are all too likely to read a p value of 0.05 or less as, "the probability that the hypothesis is *true, given the evidence of statistical significance". As many textbooks state (but almost in vain), a p value is "the probability that this evidence would arise if the null hypothesis is true" . Only Bayesian statistics yield statements about the probability of the hypothesis being true given these data which I have obtained. Once that lightbulb goes on, you are left with a convert singing the statistical equivalent of "Amazing Grace" ("...was blind, but now I see..") - hence the messianic devotion similar to a new religious convert who has finally "seen the light".*

***

I think the reason for such devotion is directly related to human nature of learning. We prefer things that are simple, flexible and common sense. Bayes theorem is precisely that. Simple because it allows us to calculate the probability that something is true, and flexible because it works even when the evidence is incomplete.

***

Possibly the strong enthusiasm is due to the fact that it allows to "get going" on problems that otherwise are "no go" zones (as per the description of Frequentist reactions in the book to certain problems) although there may be a practical need to address those problems. There's also an "all you need to do is" factor, which seems highly seductive, a hammer for all seasons, so to speak, a silver bullet. Tool bias, if you will. Two things I'd try to keep in mind myself here are Hume's problem of induction, the problem of even a more nuanced prior in the first place ("impossible/can't go wrong"), and the risk of focusing on likelihood that results from confirmation bias (Taleb's Turkey Problem - that guy who feeds me every day may end up killing me.

Week 3

Reading covered: Chapters 9-13 of The Theory That Would Not Die.

Question One

Over and over again, we see a pattern of Bayesian statistics being ignored until it proves to be practically valuable - to work in solving or predicting a real world problem. In war, in business, in computing, and so on. The theory itself is re-derived over and over by new applied thinkers. Why do you think there was such a hostility towards the methodology even as time went on, and it had proven to be so valuable? Can you think of any examples of other models that were considered "guilty until proven innocent" so repeatedly?

> *I recently read Thomas Kuhn's "The Structure of Scientific Revolutions", which was referenced in our last reading group book ("Ubiquity"). The Bayes story reminds me of Kuhn's hypothesis: new scientific discoveries are met with hostility. Tension builds when the current paradigm (in this case frequentism) can't explain certain phenomena. Eventually the pressure builds to a critical state - and adherents of the old paradigm die or retire - and the new paradigm can break through. So, if we believe Kuhn's hypothesis, maybe there's nothing unique in this story.*
>
> *The fact that Bayesian statistics has been vindicated then ignored repeatedly does seem unique, though. I'm struggling to think of a comparable situation.*

***

> *Similar line of thinking - Galileo, Copernicus ... many of the major scientific revolutions necessitated a gestation period for human inertia to catch up with change... It's the old - * this is crazy * it doesn't work * it's not interesting * of course it's like that * We actually discovered this Sequence .... Couple with financial and societal incentive, both personal and institutional to maintain the status quo and the "new" takes time to take a hold. And maybe that is a good thing from an overall societal cooperation perspective. What it really drives home also, is how much the scientific process is shaped by the deeply human, subjective, and how important rhetoric is even in science. Because once you have solved the problem on the blackboard, you have to face all that humanity out there and convince them that you are right. There's a number of occasions in the book when "good table side manner" makes all the difference in the Theory's progress. Being right also seems a question of when, and with whom.*

***

...it seems to me that the potential strength of the Bayes model lies in building an iterative process whereby input probabilities are constantly updated with every computational cycle...this approach seems to offer a robust and adaptable algorithm for asymptotic improvements in prediction...the critiques regarding single-cycle Bayesian logic should become moot with increasing iterations and precision...

Question Two

Where do you see the danger for the misuse or misapplication of Bayes' theory - in other words, in what sense might the "frequentists" -- or more generally, the Bayesian skeptics -- have a point? Is there a spot or place where Bayesian statistics might overreach, become inappropriate, or simply provide misleading answers?

Many of the problems that come to mind are not unique to Bayes, but general problems of thinking tools - essentially, looking for the keys where the light is, with the wonderful feeling that just collecting more evidence there will make it all go right. There's also a (gut feeling) question in my mind what happens when the intuitively more iterative approach of Bayes meets a problem that requires frequent qualitative leaps to get resolved. If you start on the wrong billiard table, looking for the wrong ball, you're unlikely to go far - but you may be going there for a while before you realise since you are dutifully updating... And curious what people who work practically with Bayes think. I don't.

***

I see a risk of misuse in situations that rely on subjective probabilities from surveys of experts. If the pool of "experts" aren't independent, but can influence each other, you have a biased input.

Another risk, not exclusive to Bayes theorem, is the idea that the unlikely event can't happen. i.e. I think some Bayesian models predicted Hillary Clinton winning the last U.S. Presidential election with like 80% probability and people were shocked when Trump won. Many people interpreted 80% as a sure thing, ignoring the fact that the unlikely event can still happen 1/5 of the time. Neglecting that fact can get you into trouble.

***

What immediately comes to mind as a danger is thinking that because we've run a formula, no matter how little data, there is accuracy/precision in the result. In other words, false precision. That said, I do think there's a place for subjectivity, and at least Bayes (to my knowledge) would allow one to quantify the result of using a "feel". In this week's reading there was mention of using frequency-based data in the prior - I think this is an obvious middle ground that incorporates the best of 'both worlds'.

***

The dangers might lie in the weighting of the priors being used. According to the book, Bayes' rule is a natural for decoders who have some initial guesses and must minimize the time or cost to reach a solution; So, it would be my opinion that the initial guess, if no research is available, could wrongly influence the Bayes outcome.

Week 4

Reading covered: Chapters 14-17 of The Theory That Would Not Die.

Question One

A book like this presents some key aspects and history of a mental model -- in this case, Bayesian updating -- without necessarily describing the theory in great detail. In the end, did you find it helpful resource for understanding and, importantly, _using_ the model -- if so why, and if not why not? What is your overall reaction to a book like this?

> *In my notes about this book I expressed some frustration about what I thought was padding, generic personal history small-talk items in place of insight -- who cares about the name of the boat Alan Turing took, what did he do? Then I ran across this Seneca quote, via Ryan Holiday, that seemed so very close: "We haven't time to spare to hear whether it was between Italy and Sicily that he ran into a storm or somewhere outside the world we know–when every day we're running into our own storms..."*

***

> *Candidly, this was my least favorite book we've discussed in this group. It was obviously very well researched, but came across to me as a list of anecdotes rather than providing much real insight. Some of the stories and applications of Bayes theorem were interesting and learning more about the controversy surrounding it was interesting, but I found myself losing interest about halfway through.*

***

> *I had a more positive reaction to this book than the previous replies, but perhaps that's because I am someone that routinely works with Bayes' theorem. That being said, I understand the frustrations with others. Bayes' theorem has the backbone of conditional probability, which is perhaps the most powerful tool (or mental model) in all of science: how likely is Y given X has happened. Bayes' theorem using this in a clever way allows for forecasting and answering interesting questions.*
>
> *So I found the historical tidbits interesting, especially when there was a name that came up that I recognized or a story I had heard about but got a better understanding of through its inclusion in the book. However, this historical treatment does not allow one to /use/ the mental model, but maybe this is specific to this model because it is at its roots mathematical. There are literally dozens of entry-level books on how to use and apply Bayes theorem, but each book requires a different amount of mathematical knowledge/numerical skill so they aren't going to be great for a general knowledge book club.*

That isn't to say that historical-based books are useless. I found Steven Johnson's Where Good Ideas Come From, an interesting read and its description of innovations allowed the reader to see how mental models were used to come to those innovations.

I think one thing to take away from this book not related to Bayes' theorem is the idea that there are possibly a lot of great ideas out there that may be waiting to be "rediscovered." With the advancement in computing power, especially in the last 20 years, old mathematical gimmicks are becoming powerful analytical tools because carrying out the calculations are now computationally feasible. I imagine there are great non-mathematical mental models waiting to be rediscovered too, perhaps with a tiny minority of practitioners using them to their benefit.

***

I am also in the camp of the frustrated, I'm afraid. This was my first FS read-along since Living Within Limits, but I've stayed away from the discussion, simply as I've felt pretty unable to contribute.

Part of that is down to what I felt the book failed to give me (more below) - but I also caught myself suspecting I wasn't reading/thinking deeply enough. So I'm looking forward to the summary of the discussion to try and benefit from what others have said.

From the brief notes I took (yet to process them properly), it became clear that Bayes gives us a way of dealing with uncertainty. But I just didn't get enough from the book to help me apply that. I've saved a few articles to read on how Bayes has been applied in marketing, but I just didn't understand enough from the book to get me going with that.

The Sally Clark case covered in the appendix painted a stunningly clear picture of just how revelatory a Bayesian approach could be. But despite being moderately mathematical, certainly when I was younger, I still found myself not fully mentally grasping priors/likelihood/posteriors - and how to instantly map these ideas into a real world scenario.

That left me wondering who the book was for - at points, it felt like quite a bit of maths proficiency was assumed, which left me wondering where all the maths was for readers to feast on! And certainly some of that maths could have made certain passages (e.g. on Turing's work) so much more vivid - and perhaps therefore easier to learn from.

***

Perhaps I have the wrong expectations - I thought the book would give me tools and practice to internalize this important mental model. Instead it felt like watching commercials throughout and never getting to the main course. The book did convince me I need to get this tool in my active memory and how to apply in real life but we just didn't get there, though that was the motivation in starting to read it in the first place.

Question Two

I'm going to ask you to think about applying it. McGrayne describes areas like cancer detection where simple Bayesian analyses are helpful in helping people understand what test results may (or may *not* mean for them). Can you think of some areas of your work and/or life that you feel a Bayesian approach -- even if applied generally -- could have a useful impact?

I think there are tons of areas where we can apply the ideas Bayes' theorem without actually having to do any computation. The problem is probably with our brain which doesn't like having to re-evaluate our beliefs or decisions.

We generally have a very strong set of priors when it comes to our own beliefs. For example, if you identify politically as a liberal or conservative, that is likely going to give you a strong prior about your beliefs on the efficacy of a piece of policy. If you receive new information about the efficacy of this policy (say from a scientific study or analysis of the policies performance), are you going to use this new information to update your beliefs on the policy? Maybe one study won't move your beliefs, but perhaps 5, 10, 50 of them should.

Or perhaps you decided to do some project at work because you believe that it will have some greatly positive outcome. You get new information as you're working on this project that indicates that this amazing outcome you saw from the start may not be that amazing in the end. Should you continue to work on the project or reevaluate and see if it is still worthwhile? This could be seen as another application of Bayesian reasoning.

***

There are a number of projects I can think of in which I might be able to use Bayes for an actionable outcome. The main one being expected ROI's for clients on energy saving IoT devices at their assets (smart valves for chillers, air condition motion sensors ...). Clients usually give us incomplete information (ie, they give us utility rates but can't pinpoint usage within an asset). So we have to calculate based on incomplete info. so I'd like to use Bayes on my next ROI estimation.

***

I think the medical test examples in the book can be generalised to almost any important decision in one's personal life. You can't identify every bit of uncertainty for those decisions, nor can you control the bits you do find, but by capturing prior beliefs explicitly you can make a better decision. This does not have to be explicitly mathematical.

One thing not really central to the book but essential to actually using these methods: humans suck at dealing with percentages. Natural frequencies are much easier for people to answer questions of the form "what is the probability I have disease X given result Y on test Z". They will end up using Bayes' Rule even if they don't know what that equation looks like.

***

I second the sentiment about making a distinction between the mathematical and the philosophical aspect. For those of you that may want to give Bayes a practical "mathematical" try, I am using the following approach. I am trying to work my way back from the application to the theory looking at WinBUGS/OpenBUGS, which is a well established general software but targeted to a Bayes-expert audience. I hope to fall into a feedback loop of moving back and forth between theory and practice, deep-dive and abstraction. The final goal would be to, for example, manage my expectations on some assumptions, say investing, as early as possible as new data comes in.

***

As a college basketball coach I found Bayes Theory to be helpful when thinking about a multitude of areas in coaching and quick decision making. From scouting reports, play-calling, recruiting, and practice/time management this type of mental model could be extremely useful across many areas of sports, especially in the modern age of analytics.