Statistical Paradoxes in the Courtroom

Saurabh Singal
Jan 2015

Behavioral Biases vs. Statistical Paradoxes

- The pop-behavioral finance & pop-psychology books have had such a good run that most of us know all about various kinds of biases.
- Anchoring, Frame Dependence, Mental Accounting, Availability Bias, Representativeness, Overconfidence, Self-Attribution Bias, Illusion of Control, Confirmation Bias, Hindsight Bias, Recency Bias, the Winner’s Curse, Endowment effect, Status Quo Bias, Loss Aversion, Regret Avoidance, Lottery Effect,…
- However, there are some very important **statistical biases** related to sample selection that are less commonly talked about.
- Theses biases can induce spurious correlations.

Statistical Paradoxes and Fallacies

- Statistical Paradoxes related to the Bayes' Theorem
  - Base Rate Neglect
  - Conservatism
- We will warm-up by seeing Bayes' Theorem in action before going through examples of bias in the courtroom.
  - Boy Girl Paradox
  - False Positives Paradox
- We will cover three specific examples from the Courtroom
  - Prosecutor's Fallacy,
  - Defence Attorney's Fallacy,
  - Errors in Forensic use DNA match

Biases related to Bayes' Rule

$$P(X|D) = \frac{P(X \cap D)}{P(D)} = \frac{P(D|X) \times P(X)}{P(D)}$$

Which yields

$$P(X|D) \propto P(D|X) \times P(X)$$

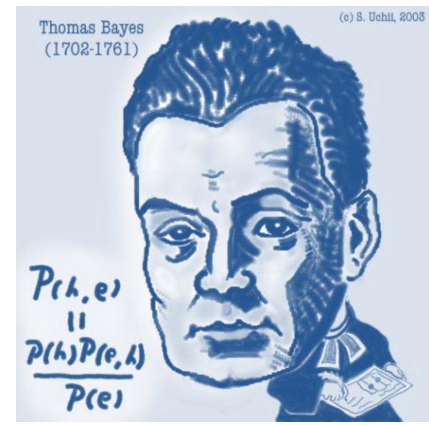

Where

$P(X|D)$ is the posterior probability

$P(D|X)$ is the likelihood

$P(X)$ is the prior probability

The three examples (Prosecutor's Fallacy, Defence Attorney's Fallacy and errors in interpreting DNA match probability) all are connected to conditional probability and Bayes' Theorem.

The Boy-Girl paradox…

- Let us start with a fun example which can be solved with or without the use of Bayes Theorem. This will give us a taste for what Bayes' Theorem can do for us!

- A new neighbor moves in next door. You learn that he has two children.

- You see that one of the children is a boy.

- What is the probability that the other child is a girl?

Boy-Girl Paradox: using Bayes' Theorem

Cases:
1. Both are girls (GG) : P(GG) = 1/4,
2. Both are boys (BB) : P(BB) = 1/4,
3. One boy, one girl- two case, (GB) and (BG)- with probability= (1/4)*2= 1/2

- Now we add the additional assumption that "at least one is a boy" : call this B.

Using Bayes' Theorem, we find:

- $P(BB|B) = P(B|BB) \cdot P(BB) / P(B)$

- $P(B|BB)$ = probability of at least one boy given both are boys = 1
- $P(B)$ = probability of at least one being a boy, which includes cases (2) and (3)
  = 1/4 + 1/2 =3/4

Finally, **P(BB|B)** = (1 * 1/4)/ (3/4) = **1/3**

So the probability other child is a girl is 2/3

Boy Girl Paradox… Intuitive Solution

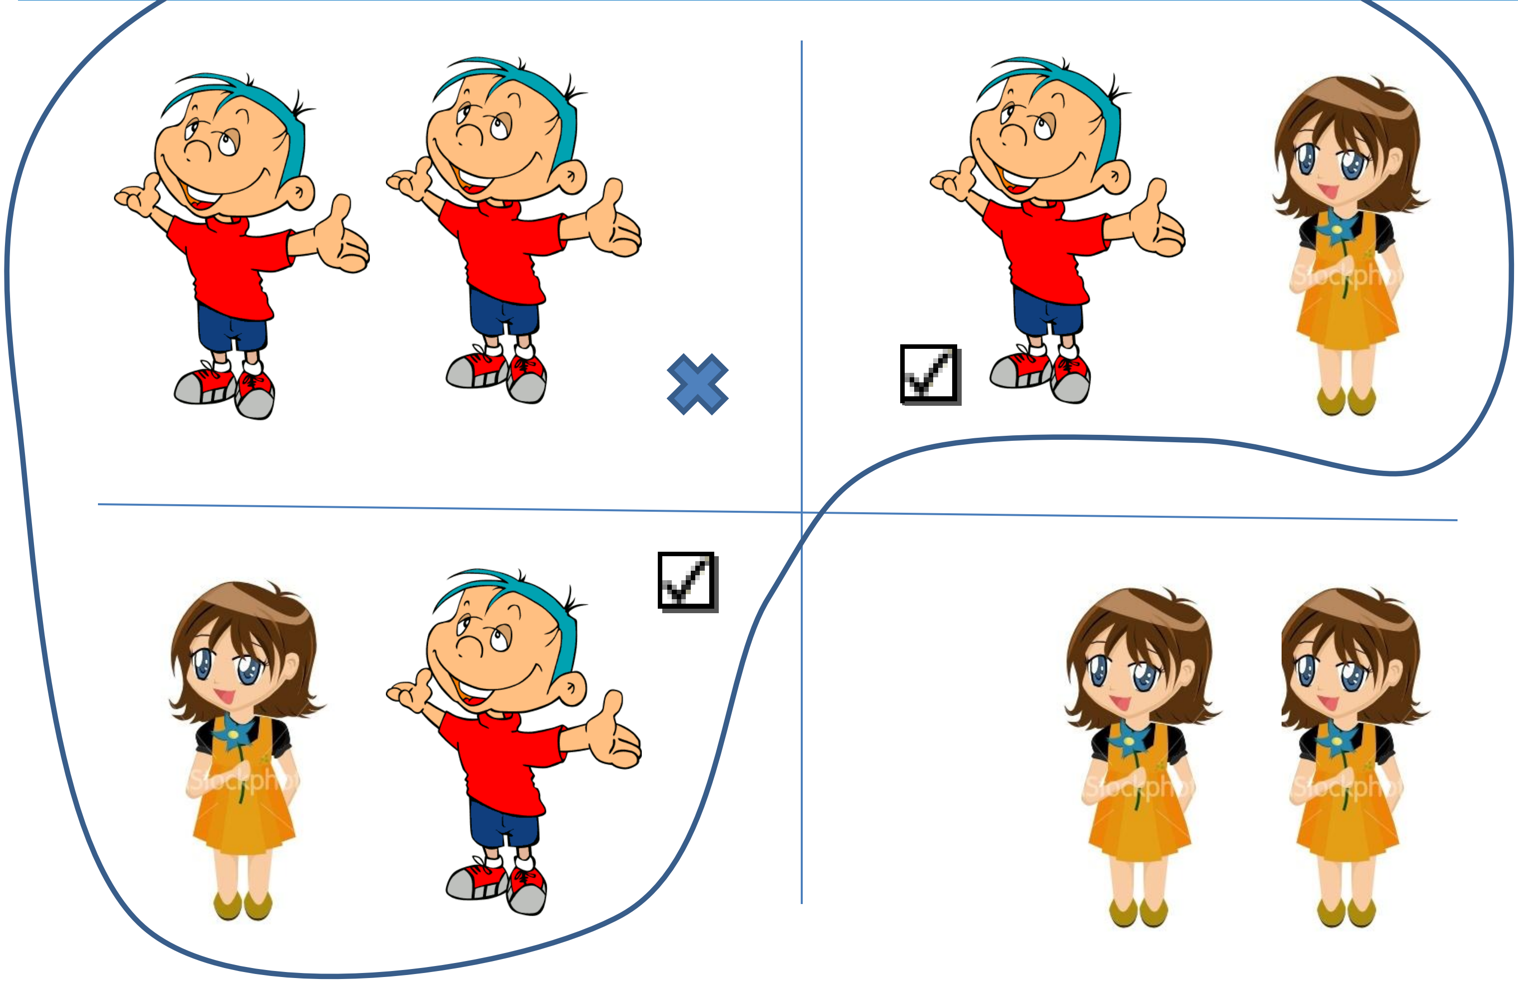

Out of 3 possible cases with at least one boy, there are two cases with a girl-

$Prob(other\ child\ is\ a\ girl) = \frac{2}{3}$

False Positive Paradox
Example from Medical Testing

This example will prepare us for the Prosecutor's Fallacy.

Testing for HIV

Here **positive** test MEANS test says "**there is an infection**"

This example will set the stage for the **Prosecutor's Fallacy.**

- About 0.01 % men are affected with HIV.
- If an ***infected*** man is tested, there is a 99.9% chance the test result is positive.
- If a man is ***not infected***, there is a 99.99% chance he will test negative.

What is the chance a man who tests positive is infected with the HIV virus?

False Positive Paradox-Answer

$$\text{Prob}(\text{Infected}\,|\,\text{Test Positive}) = \frac{\text{Prob}(\text{Infected AND Test Positve})}{\text{Prob}(\text{Test Positive})}$$

$$= \frac{\text{Prob}(\text{Test}+ve\,|\,\text{Infected}) \times \text{Prob}(\text{Infected})}{\text{Prob}(\text{Test}+ve\,|\,\text{Infected}) \times \text{Prob}(\text{Infected}) + \text{Prob}(\text{Test}+ve\,|\,\text{NOT Infected}) \times \text{Prob}(\text{NOT Infected})}$$

$$= \frac{99.9\% \times 0.01\%}{99.9\% \times 0.01\% + 0.01\% \times 99.99\%} \approx \frac{0.01\%}{0.01\% + 0.01\%} = \frac{1}{2}$$

Probability(Infected | test is positive)

= 49.98% ≈50%.

This is far lower than what most people guess.

It is easy to forget that the incidence rate of this disease is very low.

*Both the **test accuracy**, and the **incidence rat**e will influence the probability of a positive test result.*

False Positive Paradox –Answer (without tears)

Assume 10,000 men.

Of these 10,000 men, 1 is infected. He will test positive.

Of the 9,999 men, 1 tests positive (*9999 × 0.0001 = 0.9999*)

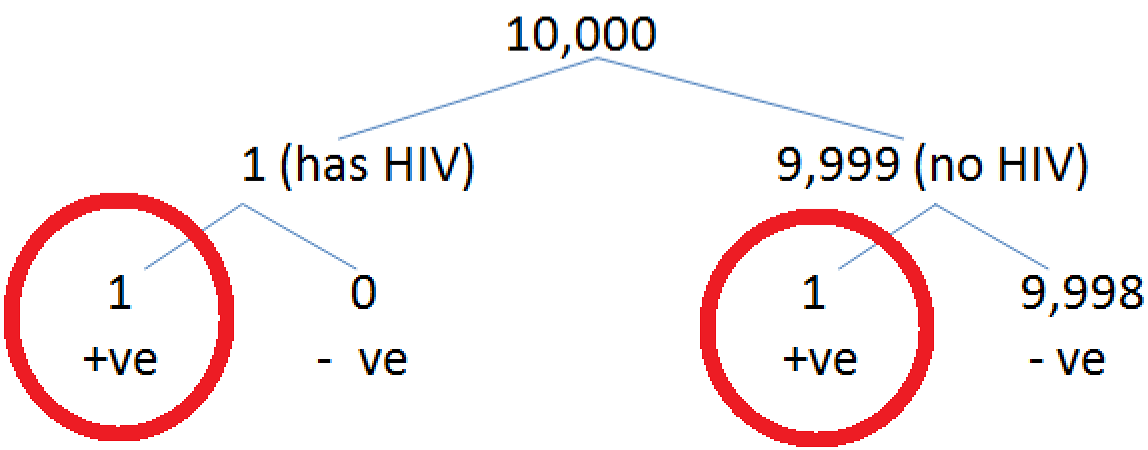

Two men test positive, of which only one has HIV.

So if he tests positive, there is only a 50% chance a man has HIV. False positive Paradox is a special case of Base Rate Fallacy.

False Positive Paradox- Conclusion

- A=The universe set, all people
  B= set of people with disease
  C= set of people testing positive
- As C becomes larger relative to B, we will encounter this paradox.

N.B. - The Figure is not to scale

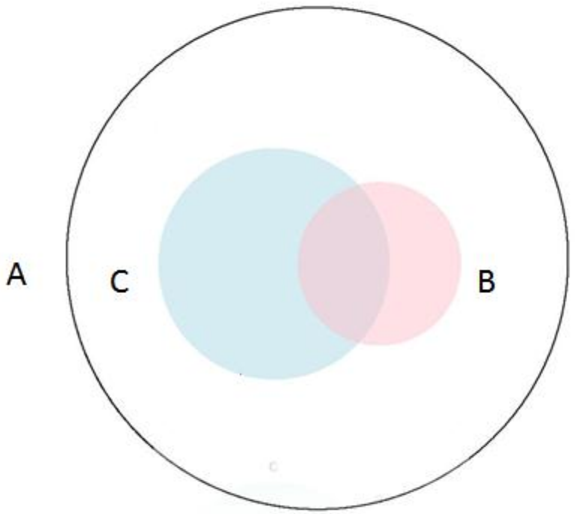

- A "false positive" is when the patient does not have the condition yet a positive result occurs (test says patient has disease).
- This paradox occurs when
  - False positives are more than true positives
  - The overall population has a low incidence of a condition
  - The incidence rate is lower than the false positive rate

Bayes' Theorem Revisited

$$P(X|D) = \frac{P(X \cap D)}{P(D)} = \frac{P(D|X) \times P(X)}{P(D)}$$

- Sometimes we can use other more visual & intuitive methods to examine such probabilities as in the False positive Paradox and the Boy Girl Paradox.

Prosecutor's Fallacy

Example from the courtroom

- **In the previous example we found**

Prob (infected| positive test) = 49.9775%
Prob (positive test| infected) = 99.9%

- **The two are different but it is easy to be confused between them.**
- **This kind of confusion leads to a Paradox called the Prosecutor's Fallacy or Juror's Fallacy.**
- **It is similar to the False Positive Paradox.**
- A Prosecutor falls victim to the Prosecutor's Fallacy when he **wrongly assumes that the prior probability of a random match is equal to the probability that the defendant is innocent**.

Prosecutor's Fallacy-2

Example from the courtroom

- **Replace "positive test" by "evidence" and "infected" by guilt in the previous example.**
- **Prob(Evidence) can mean, for example,**

the Probability that blood found at the crime scene matches the suspect's blood type .

- **Prob(Guilt| Evidence) is not the same as**

  Prob(Evidence| Guilt)
- **When a prosecutor or juror bases decisions on the evaluation of Prob(Evidence| Guilt) instead of evaluating Prob(Guilt| Evidence), this is called Prosecutor's Fallacy or Juror's Fallacy.**

Prosecutor's Fallacy-3

Example from the courtroom

- **For example, there is a rare type of blood found at the crime scene and only 1 person in a thousand has it.**
- **Further, the suspect's blood type matches this blood type.**
- **The prosecutor argues that, if the suspect did not commit the crime, the chance of match is 1 in a thousand.**

    Prob (Blood type Match |Innocence)= $\mathbf{10^{-3}}$

- **He then wrongly equates Prob (Blood type Match |Innocence) to Prob (Innocence| Blood type Match).**
- **And concludes that Prob (Guilt| Blood type Match)**

    = 1-Prob(Innocence| Blood type Match)

    = 99.9%

Case Study 1: People v. Collins

Another Example of Prosecutor's Fallacy

- On June 18, 1964, Juanita Brooks was attacked in an alley near her home in LA and her purse stolen.
- A witness saw a woman dressed in dark clothes running from the scene, who
- was blond
- had a pony tail
- fled from the scene in a yellow car, which was
  - driven by a black man who had a beard and who
    - sported a mustache

The police arrested a couple, Janet and Mark Collins, who fit the description.

Case Study 1: People v. Collins-2

Another Example of Prosecutor's Fallacy

- The prosecutor invited an instructor in mathematics to testify about the multiplication rule in probability.
- The Prosecutor then suggested the following probabilities that the jurors could use:

Evidence	**Probability**
Girl with blonde hair	1/3
Girl with ponytail	1/10
Partly yellow car	1/10
Man with mustache	1/4
Black man with beard	1/10
Interracial couple in car	1/1,000

Case Study 1: People v. Collins-3

Another Example of Prosecutor's Fallacy

- Using the probability rule for calculating joint probability, the Prosecutor multiplied these probabilities.
- The prosecutor claimed that the probability that a randomly chosen couple would have these characteristics was 1 in 12 million.
- He then said that the chance that the defendants were innocent was therefore only 1 in 12 million, i.e.,

Prob(Innocence | Match) = 1/12 million; and

Prob(Guilt | Match) = 1 – (1/12,000,000)≈ 1 or 100%.

- The jury convicted the Collinses of second-degree robbery.
- The California Supreme Court overturned this conviction on appeal

Case Study 1: People vs. Collins-4

Another Example of Prosecutor's Fallacy

The California Supreme Court overturned this conviction on appeal on several grounds

- The match probability calculation rests on certain dubious mathematical assumptions:
  - these 6 probabilities had no evidentiary foundations and were subjective estimates;
  - Probabilities cannot be multiplied together unless they are independent
- There was a logical flaw in the Prosecutor's Reasoning. *Even if the match probability is 1/12 million, the defendant could still have a high probability of being innocent*.

Case Study 1: People v. Collins-5

Another Example of Prosecutor's Fallacy

- For us, this last reason is most important as it relates to the crux of **Prosecutor's Fallacy**:

  Prob (Match) ≠ Prob (Innocent | Match)

But the Prosecutor **WRONGLY** claimed that

Prob(Innocence|Match) = Prob(Match)

and then equated

Prob(Guilt|Match) to 1- Prob(Innocence|Match)

- Suppose the Match Probability is correct, and further suppose the reference population was 24 million.
- In that case, 24,000,000 × 1/12,000,000 = 2, and we could end up with 2 matching couples.
- Then, P (Innocent |Match) = ½ = 0.5 or 50%.

Case Study 2: People v. John Puckett

Multiple Comparisons and Cold Hits on a DNA Database

- **Cold Hit-** a crime-scene sample is found to match a known profile in a database, resulting in the identification of a suspect based ***only on*** genetic evidence.
- **Random Match Probability** (RMP): the probability that the DNA profile of a randomly chosen individual matches a particular DNA profile.
- Database Match Probability or DMP is the probability that the database search had hit upon an innocent person.
- *$DMP = RMP \times N$, where N = number of people in the database*

Case Study 2: People v. John Puckett-2

Multiple Comparisons and Cold Hits on a DNA Database

- In the 1970’s, a young nurse, Diana Sylvester was found murdered in Francisco.
- The case was unsolved for decades; but in 2004 a search of California’s DNA database of 338,000 criminals found that badly deteriorated DNA from the assailant’s sperms were “matched” to a 70 year old criminal John Puckett.
- No other evidence connected the crime to John Puckett who insisted he was innocent.
- At Puckett's trial, the prosecutor told the jury that the chance of such a coincidence was 1 in 1.1 million. (Random Match Probability or RMP).
- RMP is the WRONG statistics to use- probability of guilt is not 1-RMP.
- We need to use the DMP (Database Match Probability).
- In this case, DMP= RMP × N= 338,000 × (1/1,100,000) = 30.73%
- ***Prob (database search found an incorrect match) ≈*** **1 in 3**.

Case Study 2: People v. John Puckett-2

Multiple Comparisons and Cold Hits on a DNA Database

- A fallacy can arise from **multiple testing**, such as when evidence is compared against a large database.
- The size of the database elevates the likelihood of finding a match by pure chance alone-DNA evidence is soundest when a match is found after a single directed comparison.
- The existence of matches against a large database where the test sample is of poor quality is very likely by mere chance.
- In 2004, in Chicago DNA match using skin cells from the crime scene matched a woman who was serving prison time at the time of the robbery!
- Coincidental matches are much more common than judges and juries realize, due to the sheer size of the DNA databases.

Defence Attorney's Fallacy

Yet another example from the courtroom

- In the defense fallacy, a defense attorney infers that if the Random Match Probability is p, then in a town of N people, $N \times p$ people could have committed the crime.
- E.g., a suspect in a town of 1,000,000 people has a Random Match Probability (RMP) of 1/10,000 .
- The defense attorney may imply that this suggests that any one of 100 people in the city may be the culprit.
- This fallacy underemphasizes the weight of this evidence and to ignore other evidence - not all people the other people had similar access to crime scene, motive etc.

We are all Bayesians now!

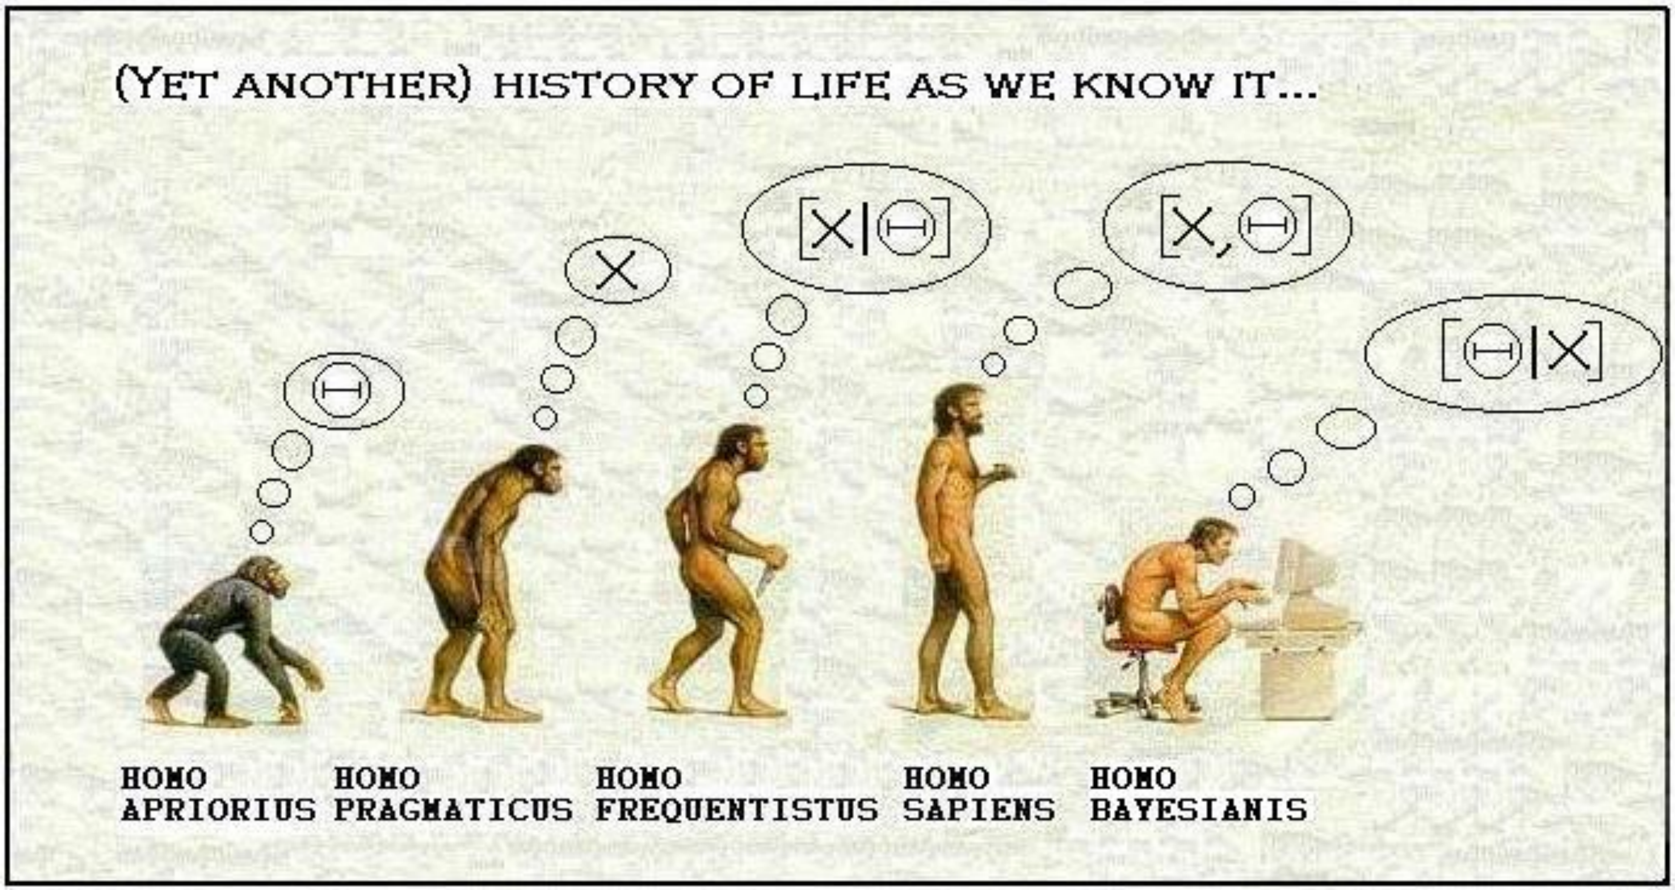

Conclusion…

Θ - Will I be attacked?
X - I heard a noise nearby

- **A priori** - Oh no! someone might attack me!!
- **Pragmatic** - What's that noise? Prepare for an attack!
- **Frequentist** - How often did I hear a noise before I was attacked?
- **Sapiens** - Thinks about how often he hears noises and how often he was attacked, but does not have the intuitive or instinctive ability to combine the ideas, as we saw in this talk….
- **Bayesian** - Answers the question, ***if I hear a noise, what is the chance of me being attacked?***

So be aware that our brains do not intuitively give Bayesian results.

You may not need a computer but if you find yourself wrestling with a problem like the one above, at least reach for a pencil and paper.