Some notes on Bayes Rule

1 Introduction

In order to properly understand the Bayes-Nash equilibrium concept one must first understand Bayes' rule. Bayes' rule is a method for updating probabilities based on new information. The most intuitive example I can give is an auction example where bidders are calling out bids (this is what people usually think of as an auction, at least until EBay became popular, but EBay operates on some of the same principles). If Bidder 1 "knows" that all bidders have values drawn from a uniform distribution over the integers from \$1 and \$100 and that all bidders seek to make a profit (so that value minus price paid is ≥ 0) then Bidder 1 knows that if Bidder 2 bids \$11 Bidder 2 must have a value of at least \$11. Thus, Bidder 1 initially assigns a $\frac{1}{100}$ probability that Bidder 2 has any of the values between \$1 and \$100. After seeing Bidder 2 bid \$11 Bidder 1 updates his beliefs about Bidder 2's value. Bidder 1 now assigns a 0 probability to Bidder 2 having a value between 1 and 10, and a $\frac{1}{90}$ probability to Bidder 2 having any particular value between \$11 and \$100.[1]

> Here is another example. Consider 2 urns. Urn A has 1 red ball and 4 white balls. Urn B has 2 red balls and 2 white balls. One of the urns is chosen at random (with probability $\frac{1}{2}$) and a ball is drawn. The ball is white. What is the probability the ball came from Urn A?

Prior to observing the ball being drawn we would say that the probability that Urn A is chosen is $p = 0.5$ because the urns are chosen at random. But now there is new information – we have observed a white ball being drawn. This observation allows us to update our belief about which urn was chosen. Let me reiterate the probability we want to find: The probability that Urn A was chosen conditional on observing a white ball being drawn.

What do we know? We know the probability with which Urn A is chosen ($\frac{1}{2}$) as well as the probability of getting a white ball from Urn A ($\frac{4}{5}$). Multiplying these two together tells us the probability of observing a white ball AND Urn A ($\frac{2}{5}$). It does NOT tell us the probability that Urn A was chosen conditional on observing a white ball, just the odds of getting Urn A and a white ball.[2] What else do we know? We know the overall probability of choosing a white ball. There is a $\frac{4}{5}$ chance of getting a white ball from Urn A and a $\frac{1}{2}$ chance of getting a white ball from Urn B. Each of these urns is chosen with probability $\frac{1}{2}$, so overall there is a $\frac{4}{5} * \frac{1}{2} + \frac{1}{2} * \frac{1}{2} = \frac{13}{20}$ chance of getting a white ball. Now, the probability of getting a white ball AND Urn A ($\frac{2}{5}$) is just the product of the probabilities that Urn A was chosen conditional on observing a white ball being drawn (this is the one we want to find) and the probability of a white ball being chosen overall ($\frac{13}{20}$). So we have $x * \frac{13}{20} = \frac{2}{5}$, where x is the probability for which we are looking, and solving for x we have $x = \frac{8}{13}$. Thus, based on seeing a white ball being drawn the probability that Urn A was chosen is now $\frac{8}{13}$ (about 61.5%). Recall that our initial guess (before observing the white ball being drawn) was $\frac{1}{2}$. This is known as Bayesian updating. We will now formalize this.

[1] Note that this assumes somewhat limited rationality on Bidder 1's part. It may be that when the game is solved from the end that only 1 particular "type" would ever have placed a bid of \$10 (maybe that is someone with a value of \$62). Thus, Bidder 1 might actually be able to infer Bidder 2's exact value from a particular bid. But let's stick with the easy example first.

[2] We can calculate the odds of getting Urn A and a red ball ($\frac{1}{10}$) as well as the odds of getting Urn B and a white ball ($\frac{1}{4}$) and the odds of getting Urn B and a red ball ($\frac{1}{4}$). Note that the sum of these 3 probabilities and $\frac{2}{5}$ is equal to 1 (one of those 4 event has to happen).

1.1 Bayes' Theorem

We first need to define an event. An event might be "a white ball is drawn" or "Urn A is chosen" – essentially, something happened. We will let "Urn A is chosen" be event A and "a white ball is drawn" be event B. We let $P(A)$ denote the probability that event A has occurred and $P(B)$ denote the probability that event B has occurred. So, our problem is to find the probability that event A occurred given that we observed event B. We say this is a conditional probability and write $P(A|B)$. Another probability that we can define is the joint occurrence of event A and event B. This is the intersection of the two events and is denoted $P(AB)$. We define the conditional probability $P(A|B)$ as:

$$P(A|B) = \frac{P(AB)}{P(B)}.$$

Think about what this states: The probability that A occurs given that B is observed is equal to the ratio of observing the events A and B together relative to the probability that the event B occurs in general. The question then is how to find $P(AB)$. We can start by writing the probability of observing B conditional on A as:

$$P(B|A) = \frac{P(AB)}{P(A)}.$$

If we rewrite this as:

$$P(AB) = P(B|A) * P(A)$$

then if we know $P(A)$ and $P(B|A)$ we know $P(AB)$. This would make the conditional probability of observing A given B to be:

$$P(A|B) = \frac{P(B|A) * P(A)}{P(B)}.$$

The equation above is Bayes' Theorem. Usually you know either $P(A|B)$ or $P(B|A)$. In our example we knew that if Urn A (event A) is chosen that the probability that a white ball would be drawn (event B) is $\frac{4}{5}$, so we knew $P(B|A)$. What we did not know was $P(A|B)$, but since we knew $P(B|A)$, $P(A)$, and $P(B)$ it was easy enough to calculate. Using the rule we see that:

$$P(\text{Urn A is chosen}|\text{ a white ball is drawn}) = \frac{\frac{4}{5} * \frac{1}{2}}{\frac{13}{20}} = \frac{8}{13}.$$

The assumption we will make in our games is that the players are all Bayesian, meaning they use Bayes' Theorem to update their probabilities.