CLASSROOM GAMES

UNDERSTANDING BAYES' RULE

Charles A. Holt and Lisa Anderson*

Abstract This paper uses the techniques of experimental economics to set up a classroom situation where students learn to make Bayesian decisions. The exercises allow students to discover for themselves a natural counting heuristic that corresponds to Bayes' rule and is much quicker to use in many situations. In the context of balls and urns, this heuristic involves adjusting ball counts to reflect prior probabilities. It provides a natural bridge between simple intuition and the mathematical formula for Bayes' rule that is presented in undergraduate courses in economic statistics, game theory, and managerial economics.

I. Introduction

Bayes' rule has always been a useful tool in the analysis of economic data. Recently, its importance in economic theory has increased as a result of the study of markets with asymmetric information or with uncertainty about distributions of wages, prices, or decisions of other players in economic games. The teaching of Bayes' rule can be frustrating for all concerned because a wide gap often appears between the near religious feelings of a Bayesian and the apathy of students.

It is useful to begin with a handful of realistic examples that indicate how simple intuition can be misleading. Of particular relevance to college students is a (true) story about a man who received a positive outcome on a first-stage test for the virus that causes AIDS. The test that was used had a 4 percent rate of false positives, and for simplicity, it is assumed that there were

* University of Virginia and The American University, repsectively. Support for this research was provided in part by the National Science Foundation (SES93-20617). We are indebted to Richard Cyert and the late Morris DeGroot for some perspectives on Bayes' Rule expressed here. We wish to thank Alan Auerbach, Sheryl Ball, Lola Berrade, Amy Hakim, Alan Krueger, Jonathan Skinner, and Timothy Taylor for helpful suggestions. The usual disclaimer applies.

virtually no false negatives.[1] The person committed suicide before follow-up examinations, presumably not realizing that the low incidence of the virus in the male population (about 1 in 250 at that time) resulted in a posterior probability of having the virus of only about 10 percent.[2] To explain this point in class, it can be useful to begin with a hypothetical representative group of 1000 people, and to ask how likely it is that a person with a positive test actually carries the virus, given an infection rate of 1 in 250 for the relevant population. On average, only 4 out of the 1000 actually have the disease, and the test locates all 4 of these true positives. However, among the 996 who do not have the disease, the test will falsely identify 4 percent as having it, which is about 40 men (.04x996 = 39.84). On average, the test identifies 44 of the 1000 men as carriers of the virus, 4 correctly and 40 incorrectly, which means that a positive first-stage test actually produces a less-than-ten percent chance of a true positive. This is a case in which knowing the intuition behind Bayes' rule can save lives. In the inaugural issue of this journal, Salop (1987) discusses other examples of Bayesian reasoning in legal contexts, which is useful given the significant proportion of economics majors who are considering careers in law. Applications of Bayes' rule in the context of financial markets and "information cascades" will be covered in a subsequent column.

Although such stories can arouse much more student interest than abstract "ball/urn" examples, our experience is that ball/urn examples can also be an effective teaching aid, especially when students participate actively in prediction games. This paper describes how to set up a classroom situation where students learn to make Bayesian decisions. The instructor should be warned that Bayes' rule itself is a difficult concept for students. Consequently, the level of difficulty of the suggested exercises is somewhat above the norm for this column.

---

[1] False negatives are difficult to identify, given time lags and repeat exposures, but the rate of false negatives on these tests is generally assumed to be very low. The 4 percent rate of false positives is roughly appropriate for early tests for the virus, although current versions are more accurate.

[2] Using Bayes' rule with a 1/250 rate of true positives and a .04 rate of false positives, the probability of actually having the virus given the positive test is calculated as the ratio: (1)(1/250) / [(1)(1/250) + (.04)(249/250)], which is about 9 percent. The 1/250 infection rate for the U.S. male population was used to produce a posterior that roughly corresponds to the 10 percent level reported in a discussion of this suicide in *The Economist* ("Think again," June 4, 1992, p. 96). In early 1996, the 1/250 infection rate would be more appropriate for the overall adult population, with the rate being about 1/800 for adult women and 1/100 for adult men in the U.S. The late Professor Morris DeGroot used to preface his discussions of Bayes' rule with a similar example, but in a pre-AIDS era the disease was either cancer or was unspecified (DeGroot, 1975, p. 61, problem 7).

The exercise involves showing students an actual draw of a colored marble from one of two hidden cups, with publicly announced proportions of marbles of two different colors. This is not dull if a student is asked to make a public prediction of which cup was used, with a small (say $.25) reward for a correct guess. It helps to ask for an explanation, and incorrect explanations can be used to motivate the pesky denominator in Bayes' rule. Moreover, the cup/marble examples allow the students to discover for themselves a natural counting heuristic that corresponds to Bayes' rule and is much quicker to use in many situations. This heuristic involves adjusting ball counts to reflect prior probabilities. It provides a natural bridge between simple intuition and the mathematical formula for Bayes' rule, as presented in undergraduate courses in economic statistics, game theory, and managerial economics.

II. Procedures

The classroom exercise takes about 30 minutes. It can be done with three marbles of one color, three of another color, a six-sided die, and two identical plastic cups of the large size common at college football games. Uniform-sized colored marbles are durable but may be difficult to find. A good substitute can be the colored bath-oil beads of a one-inch diameter that are available at body/bath shops. Each cup contains two marbles of one color and one of the other. For concreteness, let urn A indicate a cup containing two Red marbles and one White marble, and let urn B indicate a cup containing one Red and two Whites. A volunteer from the class goes behind a screen (or to the back of the class room) and throws the die. It is announced that urn A will be used if the result is 1, 2, or 3, and urn B will be used otherwise, but that students in the class will not see the actual die throw. Before beginning, it is useful to summarize the setup on the blackboard with Figure 1, where the dark circles indicate Red marbles and the white circles indicate White marbles.

Next, a single marble is drawn and shown to the class, and each student is asked to write down the probability that the draw came from urn A, along with their name and a short explanation of how they found the probability. (Although the word "probability" would not normally be used in a research experiment, there is no reason not to use it in an economics class.) The involvement of each student ensures a wide variety of answers, and the written explanations ensure that incorrect reasoning will not be conveniently forgotten. It is easiest to

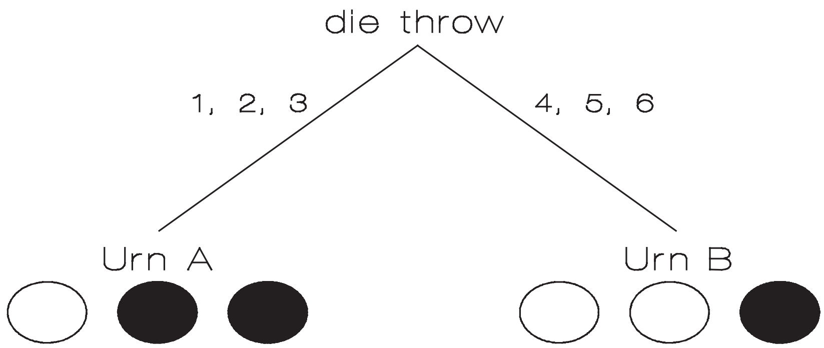

Figure 1. The Physical Setup

make copies of Figure 1 on student decision sheets that list a number of questions:

1. After the urn is selected by the throw of the die, but before any draw is seen, what is the probability that urn A is being used? _____
2. The color of the first marble drawn is ________. What is the probability that urn A is being used? _____ Explain your answer briefly:
3. The first marble was put back into the cup, and a second marble is drawn from the same cup. The color of the second marble drawn is ________. What is the probability that urn A is being used? _____ Explain your answer briefly:

Then one collects the decision sheets and tabulates the numbers of the various responses. If the first draw is a Red marble, most reported probabilities for urn A fall in the range from 1/3 to 2/3. The reported probabilities can be used to stimulate discussion and develop an intuitive version of Bayes' rule.

III. A Counting Heuristic

Suppose the first draw is a Red marble and students are asked to report the probability that urn A is being used. An answer of 1/2 is disappointingly common, even among graduate students. When you ask students why they wrote 1/2, they may say that each urn was equally likely to be selected. The response to this is: "Each urn was equally likely beforehand, but what did you learn from the draw?" If nobody chose 1/2, you can make the same point by asking someone why they did not write 1/2, "since each urn was equally likely to be selected." It is useful to begin with this relatively easy question to boost student confidence and stimulate discussion of harder questions.

Another incorrect but commonly reported probability for urn A following a Red draw is 1/3. If you ask someone who wrote a probability of 1/3 for urn A (containing 2 Reds and 1 White) to explain, you might get: "The probability of getting urn A is 1/2 and the probability of getting a Red from A is 2/3, so the probability that the observed Red draw came from A is 1/2 times 2/3, or 1/3." The best response to this is: "Then what is the probability of urn B?" Or, "The same argument requires the probability of urn B to be 1/2 times 1/3, or 1/6, but if the probability of urn A is 1/3 and the probability of B is 1/6, where does the rest of the probability go?" This leads to the realization that the probabilities must be scaled up so that they add to 1 -- that is divide the 1/3 for urn A by (1/3 + 1/6), and divide the 1/6 for B by the same thing, which yields posteriors of 2/3 for A and 1/3 for B. As before, the same point can be made even if nobody wrote a probability of 1/3 for A; simply explain the (1/2)(2/3) = 1/6 calculation and ask someone why they did not answer in this way.

Finally, find someone who wrote a probability of 2/3 for urn A, read their written explanation, and ask them to elaborate. Regardless of whether the answer was a lucky guess or a correct probability calculation, it is useful to turn everyone's attention to Figure 1 and ask if there is a simple method of counting the marbles to determine the probability that a Red draw came from urn A. With weak hints, someone will realize that all six marbles are equally likely to be drawn before the die is thrown, and that two of the three Red marbles are in urn A. It follows that the posterior probability of urn A given a Red draw is 2/3.

The calculations in the previous paragraphs are a special case of Bayes' rule with equal priors, which can also be used to update one's beliefs after receiving additional information. Suppose that the first draw of a Red marble is replaced in the cup, and the subject is told that a second draw is to be made from the <u>same</u> cup originally selected by the throw of the die. Having already seen a Red, the subject's beliefs before the second draw are that the probability that urn A is used is 2/3, or in other words, that it is twice as likely that urn A is being used. Even though the physical number of marbles in each urn has not changed, we can represent these posterior beliefs by thinking of urn A as having twice as many marbles as urn B, *with each* *marble in either urn having the same chance of being drawn*. These posterior beliefs are represented in Figure 2, where the proportions of Red and White marbles are the same as they were in urns A and B respectively. Even though the physical number of marbles is unchanged at six, the prior corresponds to a case in which the imagined marbles in Figure 2 are numbered from one to nine, with one of the nine marbles chosen randomly.

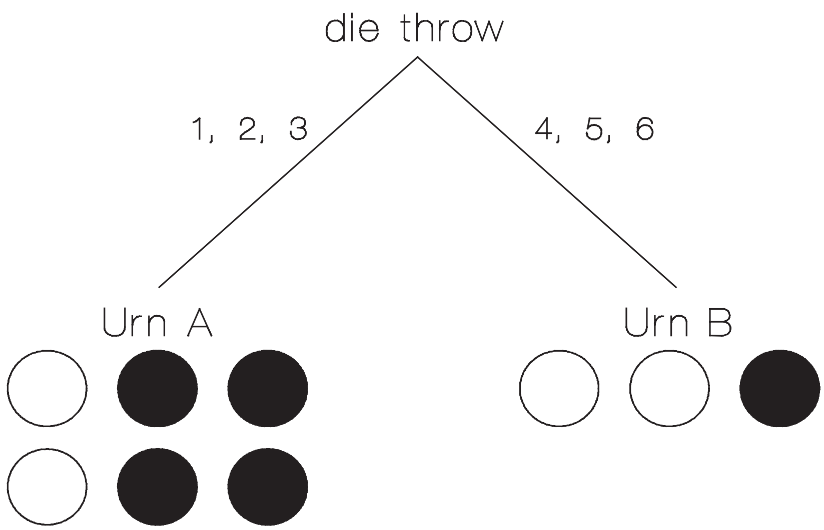

Figure 2. Representation of the Posterior After a Red Draw: Urn A is Twice as Likely

When the posterior after a Red draw is represented in Figure 2, it is clear that a White on the second draw is equally likely to have come from either urn, since each contains two White marbles. Thus the posterior after a White on the second draw is 1/2. This is consistent with intuition based on symmetry, since the prior probabilities for each urn were initially 1/2, and the draws of a Red (first) and a White (second) are balanced.

Suppose instead that the two draws were Red. As before, the posterior after the first Red draw can be represented by the urns in Figure 2, with the stipulation that each marble in either urn is equally likely to be drawn. Since four of the five Red marbles are in urn A, the posterior after seeing a (second) Red draw is 4/5. After two Red draws, urn A is, therefore, four times as likely as urn B. To represent this posterior in terms of colored marbles that are equally likely to be drawn, we need to add two more rows of three marbles under urn A in Figure 2, holding the proportions of Red and White marbles fixed.

This counting heuristic is easy to explain to students, but it is better to let them figure it out themselves as draws are made. Once they find the simple counting rule for interpreting the first draw when the urns are equally likely, as in Figure 1, make a second draw. Then ask them what the counting method would imply about the probability of urn A. Since the initial draw resulted in unequal prior probabilities (before the second draw), an incorrect answer can be used to motivate the increase in the imagined number of marbles in the urn with a higher prior probability, as in Figure 2. Regardless of whether the second draw matches the first or not, ask them to consider the other case and use the counting method to calculate the posterior for that case.

To check students' understanding, ask them what the posterior probability for urn A would be after two matching Red draws and a third White draw. Recall that two Red draws result in a posterior of 4/5 for urn A, which can be represented by listing the White/Red/Red row four times under urn A in Figure 1. The answer obtained from counting White marbles in this modified Figure 1 would be 2/3, which is intuitive since the two Red draws outweigh the one White draw, making urn A more likely. It will be necessary to give students a number of similar problems to ensure that they know how to make Bayesian calculations in this manner. It is also important to have them check the calculated posteriors against their intuition. For example, since the prior for each urn was initially 1/2, the posterior for urn A should exceed 1/2 whenever the observed number of Red draws exceeds the observed number of White draws.

III. Relating the Counting Heuristic to Bayes' Rule

This counting heuristic is quite easy for students to pick up, but it is important to relate it to the conditional probability formulas that appear in their textbooks. To make this connection, suppose there are N marbles in each urn. Let P(C|A) denote the fraction of marbles in urn A that are of color C, where C is either Red or White. Similarly, P(C|B) is the fraction of marbles in urn B that are of color C. Then there are a total of P(C|A)N marbles of color C in urn A, and there are P(C|B)N marbles of color C in urn B. If each urn is equally likely to be selected, then each of the 2N marbles in the two urns is equally likely to be drawn. Suppose the marble drawn is of color C. The posterior probability that a marble of color C was drawn from urn A, denoted

P(A|C), is just the ratio of the number of color C marbles in urn A to the total number of marbles of this color:

$$P(A|C) = \frac{number\ of\ color\ C\ marbles\ in\ urn\ A}{number\ of\ color\ C\ marbles\ in\ both\ urns}$$

$$P(A|C) = \frac{P(C|A)N}{P(C|A)N + P(C|B)N} \qquad (with\ N\ marbles\ in\ each\ urn).$$

It is worth emphasizing that this formula is only valid for the case of equal prior probabilities and equal numbers of marbles in each urn, so that the N cancels from each term on the right. In particular, divide the numerator and the denominator of this equation by 2N, which yields a formula for calculating the posterior probability of urn A when the prior probabilities are 1/2:

$$P(A|C) = \frac{P(C|A)[1/2]}{P(C|A)[1/2] + P(C|B)[1/2]} \qquad (with\ priors = 1/2).$$

A person who has seen one or more draws may not have prior probabilities of 1/2, so this formula must be generalized. This involves replacing the [1/2] terms on the right side of the with the new prior probabilities, denoted P(A) and P(B). This is Bayes' rule.

$$P(A|C) = \frac{P(C|A)P(A)}{P(C|A)P(A) + P(C|B)P(B)} \qquad (Bayes'\ rule).$$

For the previous example with equal priors and one Red draw, P(A) = 2/3, P(Red|A) = 2/3, and P(Red|B) = 1/3, so the posterior following a Red draw, P(A|Red), is 4/5.

Having arrived at the Bayes' rule formula (for two events) as it appears in the textbooks, it is important to point out that the argument was intuitive but not rigorous. It helps, therefore, to relate the general formula for Bayes' rule back to the counting heuristic. This can be done by dividing numerator and denominator of Bayes' rule by P(B)/N, to obtain:

$$P(A|C) = \frac{P(C|A)\frac{P(A)N}{P(B)}}{P(C|A)\frac{P(A)N}{P(B)} + P(C|B)N},$$

or equivalently,

$$P(A|C) = \frac{P(C|A)\,M}{P(C|A)\,M + P(C|B)\,N},$$

where $M = [P(A)/P(B)]N$. If there are N marbles in each urn, these equations show that the basic Bayes' rule formula is equivalent to imagining that the N marbles in urn A are increased or decreased to a number M, which is the product of N and the prior odds ratio: P(A)/P(B).[3]

To summarize, if there is a prior probability of 1/2 that each urn is used and if the urns contain equal numbers of colored marbles, then the posterior probabilities can be calculated as ratios of numbers of marbles of the color drawn, as in the first equation given. If the marble drawn is of color C, then the posterior that the draw was from urn A is the number of color C marbles in urn A divided by the total number of color C marbles in both urns. When the prior probabilities or numbers of marbles in the cups are unequal, then the 1/2 terms in the second display equation are replaced by the prior probabilities, as in Bayes' rule. Finally, the Bayes' rule formula is equivalent to imagining that the number N of marbles in urn A is changed by a factor P(A)/P(B), holding the proportion of color C marbles fixed at P(C|A), with the counting heuristic that each of the actual and imagined marbles is equally likely to be chosen.[4]

This discussion allows students to see how Bayes' rule can be used to determine the relative importance of prior and sample information. This exercise develops a simple counting heuristic that permits one to make Bayesian decisions without looking explicitly at ratios of conditional probability calculations. The heuristic also provides a simple explanation of the algebra of Bayes' rule.

IV. Further Reading

Instructors should be warned that students often make large errors in calculating probabilities when there are no financial incentives. For the setup in Figure 1, the draw of a Red marble produces a posterior of 2/3, but it is not uncommon to see reported probabilities that are less than 1/2. This "noise" is good for generating classroom discussions. However, economists

---

[3] When the prior odds ratio times N is not an integer, one can count fractional marbles, or increase N for both urns to a level that yields an integer-valued product of N and the odds ratio. Also, this counting heuristic can be generalized to the case of more than two events (urns) in the obvious manner.

[4] There are a number of other heuristics that are commonly used in teaching Bayes' rule. Perhaps the most common is to multiply the prior odds ratio, P(A)/P(B), by the likelihood ratio, P(C|A)/P(C|B), to get the posterior odds ratio, P(A|C)/P(B|C).

frequently care about predictions made under real economic incentives, and research on learning typically provides subjects with financial incentives. In particular, subjects are put into a money-making situation where it is in their own best interest to report a probability that corresponds to their own personal beliefs about the likelihood of some event. These methods are introduced in Davis and Holt (1993, chapter 8).

Laboratory tests of Bayes' Rule are surveyed more completely by Camerer (1995), who also discusses the related psychology literature on biases. The use of financial incentives may reduce biases, but they are not eliminated. For example, a sample of two Reds and a White is "representative" of urn A in Figure 1, in the sense that the sample proportions match the urn proportions. If a subject is asked to guess which urn is more likely in this case, the Bayesian prediction of urn A corresponds to the prediction based on a representativeness heuristic. It is possible to design other situations in which the Bayesian decision differs from one based on representativeness: i.e., if one of the urns has such a low prior probability of being used that it is still the less likely urn even after seeing a sample that exactly "represents" its proportions of colored marbles. Deviations from Bayesian predictions are more common in such situations (Grether, 1980, 1992).

References

Camerer, Colin F., "Individual Decision Making." In Kagel, John, and Alvin Roth, eds., *Handbook of Experimental Economics*. Princeton: Princeton University Press, 1995.

Davis, Douglas D., and Charles A. Holt, *Experimental Economics*. Princeton: Princeton University Press, 1993.

DeGroot, Morris H., *Probability and Statistics*. Reading, Mass.: Addison-Wesley, 1975.

Grether, David M., "Bayes' Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, March 1980, *95*, 537-57.

Grether, David M., "Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence," *Journal of Economic Behavior and Organization*, January 1992, *17*, 31-57.

Salop, Steven C., "Evaluating Uncertain Evidence With Sir Thomas Bayes: A Note for Teachers," *Journal of Economic Perspectives*, Summer 1987, *1*, 155-60.

"Think again," *The Economist*, June 4, 1992, 96.