Cognition, 11 (1982) 109–122

Some teasers concerning conditional probabilities

MAYA BAR-HILLEL*
RUMA FALK
Hebrew University of Jerusalem

Abstract

A family of notorious teasers in probability is discussed. All ask for the probability that the objects of a certain pair both have some property when information exists that at least one of them does. These problems should be solved using conditional probabilities, but cause difficulties in characterizing the conditioning event appropriately. In particular, they highlight the importance of determining the way *information is being obtained. A probability space for modeling verbal problems should allow for the representation of the given outcome* and *the statistical experiment which yielded it. The paper gives some psychological reasons for the tricky nature of these problems, and some practical tips for handling them.*

1. Introduction

A few years ago, one of the authors was visited by two friends, both professors of mathematics at the Hebrew University of Jerusalem. While discussing her experience in teaching an Introductory Probability course, she remarked that some problems, even at that elementary level, seem extremely tricky and difficult to conceptualize correctly. Skeptically, her friends asked for an example. She provided the following:

Problem 1

Mr. Smith is the father of two. We meet him walking along the street with a young boy whom he proudly introduces as his son. What is the probability that Mr. Smith's other child is also a boy? (Falk, 1978).

*This work was supported by the Advanced Research Projects Agency of the Department of Defense under Contract N00014-79-312-0730 under subcontract from Decisions and Designs, Inc. to Perceptronics, Inc. We wish to thank Baruch Fischhoff and Sarah Lichtenstein for comments on an earlier draft. Reprint requests should be sent to Maya Bar-Hillel, The Hebrew University of Jerusalem, Jerusalem, Israel.

In chorus the two mathematicians exclaimed: "Well, what's the problem? The answer is obviously____." Except here, one inserted "one-third" and the other inserted "one-half".

Problem 1 is simple enough to be presented to students at about the third session in an introductory course, shortly after introducing the concepts of event-independence and conditional probability. Two possible ways of approaching the problem immediately come to mind. One states that since (to a very close approximation) the two sexes are equiprobable, and the sexes of any two children (born separately) are independent, knowing that one of Smith's children is a boy does not affect our probability that the other is a boy, which was and still is one-half. Formally:

P (some child is a boy | some other child is a boy) = 1/2.

According to the second argument, before Mr. Smith identifies the boy as his son, we know only that he is either the father of two boys, BB, or of two girls, GG, or of one of each in either birth order, i.e., BG or GB. Assuming again independence and equiprobability, we begin with a probability of 1/4 that Smith is the father of two boys. Discovering that he has at least one boy rules out the event GG. Since the remaining three events were equiprobable, we obtain a probability of 1/3 for BB. Formally:

P (two boys | at least one boy) = 1/3.

Obviously, at least one of these two lines of reasoning must be fallacious. Consider now a variant of Problem 1:

Problem 2

We meet Mr Smith (whom we know to be the father of two) in the street with a boy. This time, he is more elaborate in his introduction, presenting the boy as his eldest child. What is the probability that Mr. Smith's other child is also a boy? (Falk, 1978).

Here, both approaches—considering the independence of children's sex and enumerating the possibilities remaining after the obtained information is registered—lead to an answer of 1/2.

The convergence of solutions to this second problem need not, however, force your position on Problem 1. For one could maintain that as information is added, it is appropriate that the requested probability changes. On the other hand, one might favor the position that since the probability that Mr. Smith has two sons would clearly have been the same had we found out that the child was in fact the *youngest*, it shouldn't affect the requested probability. The question seems to boil down to whether the additional information concerning the birth order of the accompanying son is or isn't relevant to the event of Mr. Smith's having two sons.

2. A related puzzle

A similar puzzle, formulated in terms of a card problem, has been making the rounds for some decades (see Gridgeman, 1967, for a historical survey).

Problem 3

A deck of four cards consists of the ace of spades, the ace of clubs, the deuce of spades, and the deuce of clubs. A hand of two cards is randomly dealt from this deck. What is the probability that it contains both aces if we know it contains at least one?

The answer to this problem is traditionally agreed upon to be 1/5, following the reasoning that five equiprobable hands are compatible with the conditioning event (only the double deuce hand is ruled out), and just one of these contains both aces.

Now compare Problem 3 to the following:

Problem 4

Like Problem 3, but the question is: What is the probability that the hand contains both aces if we know it contains the ace of spades?

By the traditional reasoning, since three of the six initially possible hands are eliminated, the answer is 1/3. The conjunction of Problems 3 and 4 poses a puzzle, however, since it seems that irrelevant information has nevertheless affected the probability that the hand has another ace. The information is deemed irrelevant, since the answer to Problem 4 clearly does not depend on whether the stated suit of the known ace is spades or clubs (see Gamow and Stern, 1958; Gardner, 1959).

The similarity of the Second-Ace problems to the Second-Son problems should be apparent. In fact, the two pairs of problems are nearly isomorphic, with the two dichotomous variables of sex and birth order being replaced by card value and suit, respectively. The difference is that *any* hand which includes two of the four cards is feasible, but not so any type of family of two. You can't have a family in which the first-born is a male, while simultaneously the first-born is a female. That is why the sample space contains six points in the card problems, and only four in the family problems.

3. Analysis of the problems

Information and how it is obtained

In order to gain a better grip on Problem 1, let us consider its story again. How exactly did we obtain the information that Mr. Smith has 'at least one boy'? What were the specific conditions of the 'statistical experiment' which gave rise to that datum? Recall that we met Mr. Smith walking in the street in the company of a son of his. In order to proceed, one has to spell out some assumptions about the 'real world' which would tie our possible observations to the kind of family which Smith might have. In the absence of additional information, it is natural to assume that, when setting out on a walk with one of his two children, Mr. Smith selects the child *at random*.

Let us denote the datum obtained in this meeting with Mr. Smith by B_m (we *m*et a *B*oy). Table 1 presents the probabilities of all combinations of family type crossed with the sex of the accompanying child. It can easily be seen that:

$$P(BB|B_m) = \frac{1/4}{1/2} = 1/2.$$

Likewise,

$$P(BG|B_m) = P(GB|B_m) = 1/4, \text{ and } P(GG|B_m) = 0.$$

Contrary to the second approach to Problem 1, which viewed the three remaining family types as equiprobable, they are seen not to be. Realizing that a father of two boys is more likely to pick a boy for a walking companion (in fact, it is a certainty) than is a father of a boy and a girl (in which case it is a toss-up), it becomes clear that the observation B_m renders the event BB more probable than either BG or GB.[1]

It is essential to notice that the conditioning event should be phrased *not* as 'Mr. Smith has at least one son', but rather as 'a randomly encountered child of Mr. Smith is a son'. Under the usual assumptions, the former has a probability of 3/4, the latter of 1/2.

Table 2 gives the probabilities of all combinations of family type crossed with sex and birth order of accompanying child. Note that by summing the probabilities in the two upper lines of Table 2 column by column, one gets

[1] This becomes even more apparent by carrying the case to the extreme. Suppose that Smith is known to be either the father of ten boys, or of one boy and nine girls. We meet him in the street in the company of a son. We should be more confident that he has sons at home rather than daughters, since the ten-boy family is ten times as likely to yield our observation as the nine-girl and one-boy family.

Table 1. *Probabilities of combinations of family type with sex of companion (Problem 1)*

Sex of accompanying child	Smith's family type BB	BG	GB	GG	Total
B_m = *meet* a *Boy*	$\frac{1}{4} \times 1 = \frac{1}{4}$[a]	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{2}$
G_m = *meet* a *Girl*	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times 1 = \frac{1}{4}$	$\frac{1}{2}$
Total	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	1

[a]Each cell entry gives the probability of the conjunction of the column event with the row event. Thus, $P(BB \cap B_m) = P(BB) \cdot P(B_m | BB) = \frac{1}{4} \times 1 = \frac{1}{4}$.

Table 2. *Probabilities of combinations of family types with sex and birth order of companion (Problem 2)*

Sex and birth order of accompanying child	Smith's family type BB	BG	GB	GG	Total
B_m^1 = *meet 1*st-born *Boy*	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4}$
B_m^2 = *meet* 2nd-born *Boy*	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4}$
G_m^1 = *meet 1*st-born *Girl*	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4}$
G_m^2 = *meet* 2nd-born *Girl*	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4} \times 0 = 0$	$\frac{1}{4} \times \frac{1}{2} = \frac{1}{8}$	$\frac{1}{4}$
Total	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	1

the probabilities in the upper line of Table 1. Likewise, summing the numbers in the third and fourth lines of Table 2 yields the second line of Table 1. The solution to Problem 2 may now be confirmed by reference to Table 2:

$P(\mathrm{BB}|\mathrm{B_m}^1) = \frac{1/8}{1/4} = 1/2$, where $\mathrm{B_m}^1$ denotes: we *m*et a *1*st born *B*oy.

Thus, although the information obtained in Problem 1 left three candidates for Smith's family type (BB, BG, and GB), whereas that of Problem 2 left only two (BB and BG), the posterior probability of the event BB is the same, 1/2, in both. Our analysis seems to suggest that this results from the fact that the two events, BB and BG, remain equiprobable in Problem 2, but not in Problem 1. Although this is a convenient explanation, it ignores the fact that the sample space {BB, BG, GB, GG} is not really suitable for modeling our problems, in a sense that will be elaborated in the following subsection. For the time being, we can rely on the 8-point (16-point) sample space in Table 1 (2), which was derived from {BB, BG, GB, GG} by partitioning it more finely.

Can one think of reasonable background assumptions for Problem 1 under which the events BB, BG, and GB would remain equiprobable? Imagine, for example, that Mr. Smith belonged to a culture that places very high value on having male offspring. In this culture, boys are *invariably* chosen over girls as walking companions. Under this assumption, a father of two boys is no more likely to pick a son to walk with than is one who selects (no longer randomly!) from one boy and one girl. In this culture, the observation $\mathrm{B_m}$ does not render BB more probable than BG or GB.

Another version of the Second-Son problem is presented by Gardner (1959, p. 51): 'Mr Smith says, "I have two children and at least one of them is a boy". What is the probability that the other child is a boy?' Gardner derives an answer of 1/3, based on the equiprobability of BB, BG, and GB. He is not, however, in disagreement with our analysis. He is merely addressing a different problem! Having Mr. Smith tell us that he has a son, denoted $\mathrm{B_s}$, is altogether different than discovering that fact on the basis of observing *only one* of his children. For Mr. Smith, unlike the reader, presumably is aware of the sex of *both* his children when making this statement. If, for example, Mr. Smith has one boy and one girl, it is necessarily true that he has at least one boy (Gardner's version), but it is not necessarily true that we will meet him accompanied by a male child (our version). Formally speaking, no event in the sample space {BB, BG, GB, GG} corresponds to $\mathrm{B_m}$, whereas $\mathrm{B_s}$ corresponds to {BB, BG, GB}. Gardner's version would become formally equivalent to Problem 1 only if we adopt some assumption like the 'male oriented' culture above.

Let us now apply this analysis to Problems 3 and 4. Note that although both these problems state that we *know* something, they give no clue as to how this knowledge was obtained, a crucial ingredient for selecting the appropriate model. Without it, the problem's formulation leaves the precise interpretation underdetermined. In an elegant exposition, Freund (1965) attempted to fill this gap by describing two alternative procedures whereby such knowledge could come to be acquired. He conjured a hypothetical spy, operating in one of the two following methods:

Case I

> The spy looks at our opponent's *entire hand*; he reports whether or not he sees (at least) one ace [Problem 3], or he reports the suit (*flipping a coin to decide whether to report spades or clubs* when he sees both aces in our opponent's hand) [Problem 4].

Case II

> The spy has the chance to see *only one card (randomly selected from our opponent's hand)* and he either reports whether or not it is an ace [Problem 3], or he also reports the suit [Problem 4] (p. 29, emphasis added).

By applying Bayes' Theorem, or by constructing the appropriate two-dimensional probability distributions, the reader can verify the following results. Case I: When the suit is not reported (Problem 3), the five remaining hands stay equiprobable, so the solution is 1/5. When the suit is reported (Problem 4), the three remaining hands are no longer equiprobable. However, since the posterior probability of the hand with both aces is only half that of either of the other two remaining hands (which contain a deuce along with the ace of spades), the solution is still 1/5. Case II: Here, whether the posterior probabilities remain equiprobable, (Problem 4) or not (Problem 3), the probability of two aces is 1/3. In either case, the same answer is obtained to both problems, and 'the paradox has vanished' (Freund, 1965). Other scenarios are possible, including one that would yield an answer of 1/5 to Problem 3 and of 1/3 to Problem 4. Some scenario, in any case, is called for, since otherwise, the story is incomplete.[2]

Recall the isomorphism of Problems 1 and 2 to Problems 3 and 4. Since we adopted the assumption that Smith selected the child in his company at random, the analogy is to the Case II interpretation of Problems 3 and 4. The assumption that a male is selected whenever possible would render Problems 1 and 2 analogous to the Case I interpretation of Problems 3 and 4.

[2] Freund's paper prompted a series of over half a dozen letters which were published in *The American Statistician* in a space of two years following his own. They offer a wide variety of solutions and approaches. An analysis similar to Freund's can be found in Betteley (1979).

Information and its relevance

So far, we have solved Problems 1–4 by carefully considering the statistical experiment underlying them. We have not yet dealt with another puzzling aspect of these problems, concerning the relevance of adding further specific information. In particular, should information about the birth order of Mr. Smith's son[3] affect the required probability? On the other hand, it seems that it should, since it clearly affects the number of possibilities that are eliminated. On the other hand, it seems that it shouldn't, since the answer is independent of the stated birth order. If the probability of an event is the same when conditioned on two complementary events—as 'eldest' and 'youngest' seem to be—then it must also equal the unconditioned (i.e., the total) probability, since $P(\mathrm{A}) = P(\mathrm{A}|\mathrm{B})\cdot P(\mathrm{B}) + P(\mathrm{A}|\bar{\mathrm{B}})\cdot P(\bar{\mathrm{B}})$.

To examine this issue, let us return to Gardner's 'at least one of them is a boy' version. After offering 1/3 as the probability that Smith has two boys, Gardner goes on to say: '[H]ad Smith said that his *oldest* ... child is a boy, then the situation is entirely different. Now the combinations are restricted to BB and BG, and the probability that the other child is male jumps to 1/2' (p. 51). Indeed, the additional information about the son's birth order, by affecting the number of remaining possibilities, but not their relative odds, affects the required probability. In our version (Problems 1 and 2), however, as soon as we observe Mr. Smith with a son, 'the probability that the other child is male jumps to 1/2' (Gardner, 1959). Moreover, it jumps no further when we learn the birth order of that son. Why it that? How does the probability in Problem 1 jump to 1/2 without Mr. Smith supplying us with any particular information about his son?

Let us analyze the role that birth order plays in this kind of problem. The intuitive sample space for Problems 1 and 2 consists of the elementary events {2 boys, 2 girls, 1 girl and 1 boy}. Unfortunately, this space is not uniform. Introducing birth order partitions '1 girl and 1 boy' into two equiprobable events, transforming the sample space into the uniform {BB, GG, BG, GB}. The success of the birth order variable in achieving this end hinges on its independence of the sex variable. But other distinguishing variables that are independent of sex could be employed for the same purpose. For example, one could distinguish between families in which the darker child is male and the fairer is female *versus* families in which the fairer child is male and the darker is female. When no natural variable is present, one may wish to impose one. One such, for example, is the designa-

[3]From here on, we will no longer use the double reference to the Second-Ace problems as well as the Second-Son problems. The readers are encouraged to draw the analogies themselves.

tion of dice by color (e.g., 'blue' and 'red') in problems involving more than one die. In yet other cases, a problem is formulated in terms of some such variable, but its role is implicit, or even disguised. Problem 1 is a case in point.

In dealing with this problem heretofore, we have resorted to the (birth) ordered pairs of sexes. We were aware that the conditioning event B_m was unexpressible within this sample space. Yet, we were surprised that the mere encounter with Mr. Smith's son, though unidentified by birth order, achieved the same effect on the probability of BB that Gardner's version achieved only when birth order was supplied. By labeling the child we met as his oldest, Mr. Smith distinguishes him from the other one. But why need Mr. Smith supply us with a label for distinguishing the child we met from the child at home? We can supply the label ourselves, namely 'the child we met'! When the unaccompanied Mr. Smith says "at least one of my two children is a boy", we have no way of telling *which* one he is talking about —at least until he adds that it is his oldest. Whereas, when we *meet* a son of Smith's, we know not only that he has at least one boy, but also precisely which one of them (at least) is a boy—to wit, the one we met.[4]

This suggests that we model Problems 1 and 2 using the following probability space. Denote the child whom Smith left at *h*ome by a subscript h, and the child we *m*eet with him by a subscript m. This yields a uniform sample space for describing Smith's family structure:

$$\{B_m B_h, B_m G_h, G_m B_h, G_m G_h\}$$

Upon encountering Mr. Smith in the company of a male offspring, $G_m G_h$ and $G_m B_h$ are ruled out. This is tantamount to saying that the conditioning event B_m equals the subset $\{B_m G_h, B_m B_h\}$. In this case, the events in the conditional sample space preserve their equiprobability, and so the revised probability that Mr. Smith has two sons is 1/2 (see also Jeffrey, 1968). It is now easy to see why adding the information that the son we meet is first-born increases the probability of BB no further. It is irrelevant *in the terms in which the sample space is construed*. Mr. Smith might as well have said: "This is my son, and his name is Jim". While it is hard to shake ourselves loose of the powerful habit of employing birth order for analyzing such problems, in this case birth order is but a red herring.[5]

---

[4] This reminds us of the story about the man who could not tell his two horses apart. After a series of failures to mark one of them in a manner that would set it apart from the other, he finally decided to sell the black horse and keep the white one.

[5] This is not to deny that the habit has a lot to commend it. Birth order enjoys the distinct advantages of being discrete, salient, compatible with the sequential flow of speech and writing, etc.

4. The prisoner's paradox

Another classic puzzler in probability theory (not to be confused with the Prisoner's Dilemma of Game Theory) is closely related to Problems 1–4 (see, for example, Beckenbach, 1970; Mosteller, 1965).

Problem 5

Tom, Dick and Harry are jailed in separate cells. Early the next morning, one of them will hang and two will be set free. A lottery has already determined the unfortunate one, but the night guard is not allowed to inform any prisoner of his fate. Dick can't sleep. He puts his chances for hanging at 1/3, too big for comfort. If he could just obtain some more information! Dick manages to convince the guard that by naming either Tom or Harry as one of those to be freed, he will not be violating his instructions. Dick bribes the guard into indulging him, and the guard names Harry. What is the current probability that Dick will hang?

Since at this point, only Tom and himself are still candidates for hanging, Dick judges his probability of hanging to have increased to 1/2. Suppose, however, that the guard had named Tom as the lucky prisoner. By the same token, this piece of information would also have increased Dick's probability of hanging to 1/2.

It looks like whatever information the guard discloses affects Dick adversely. Indeed, the very mental exercise of *imagining* the exchange with the guard appears to increase his probability of dying! What is going on?

To extricate Dick from this diabolical mess, let us apply the lesson gained in the previous sections. The agreement that Dick strikes with the guard implicitly assumes that the guard has no bias for naming either Tom or Harry, i.e., that if both are to be set free, they have an equal chance of being named.

Let T, D, and H denote the respective events that Tom, Dick, or Harry will be hanged, and let $\overline{H}_g$ and $\overline{T}_g$ be the guard's designating Harry or Tom, respectively, as one who will be freed (i.e., not hanged). Then,

$$P(\overline{H}_g|D) = 1/2;\ P(\overline{H}_g|H) = 0;\ P(\overline{H}_g|T) = 1.$$

In other words, the event that Harry will be named is not independent of the possibilities D and T. Employing these probabilities in a Bayesian computation will not, therefore, yield equal posterior probabilities for D and for T:

$$P(\mathrm{D}|\overline{\mathrm{H}}_g) = \frac{P(\overline{\mathrm{H}}_g|\mathrm{D})P(\mathrm{D})}{P(\overline{\mathrm{H}}_g|\mathrm{D})\mathrm{P}(\mathrm{D}) + P(\overline{\mathrm{H}}_g|\mathrm{H})\mathrm{P}(\mathrm{H}) + P(\overline{\mathrm{H}}_g|\mathrm{T})\mathrm{P}(\mathrm{T})}$$
$$= \frac{(1/2)(1/3)}{(1/2)(1/3) + 0(1/3) + 1(1/3)} = 1/3.$$

By a similar computation,

$P(\mathrm{T}|\overline{\mathrm{H}}_g) = 2/3$, and $P(\mathrm{H}|\overline{\mathrm{H}}_g) = 0$.

So Dick's probability of hanging really is unaffected by the guard's naming of Harry. The key to this conclusion lies in the realization that the conditioning event $\overline{\mathrm{H}}_g$ is defined not merely by the information obtained (i.e., $\overline{\mathrm{H}}$), but by the way in which it was obtained as well.

Here, too, it is possible to construct a statistical experiment which would seem to yield the same information, and yet increase Dick's probability of hanging. Suppose the guard has agreed to the following procedure: He will toss a coin. If Heads, he will report Harry's fate, whatever it is. If Tails, he will report Tom's fate. If Dick now hears 'Harry will be freed', his probability of hanging does change to 1/2.

5. The three-card problem

We conclude our presentation of probabilistic teasers with one last problem. This well-known problem appears in many variations (see, for example, Frauenthal and Saaty, 1979; Gamow and Stern, 1958). Although here the statistical experiment is explicit and transparent, the problem remains an intuitively misleading one.

Problem 6

Three cards are in a hat. One is red on both sides, denoted RR. One is white on both sides, denoted WW. One is red on one side and white on the other, denoted RW. We draw one card blindly and put it on the table. It shows a *R*ed face *u*p, denoted $\mathrm{R_u}$. What is the probability that the hidden side is also red?

We presented this problem to 53 Psychology freshmen taking one of our Introductory Probability courses. Thirty-five of them (66%) gave an answer of 1/2, apparently reasoning as follows: The card is definitely not WW, so it is either RR or RW. Since it was drawn randomly, these cards are equiprobable.

By now, the reader probably senses the fallaciousness of attributing posterior equiprobability to the remaining events. Clearly, an all-red card is twice as likely to show a red face up as a card that only has one red side. Hence, by Bayes' Rule, $P(\mathrm{RR}|\mathrm{R_u}) = 2/3$. Only three of our respondents gave this answer.

6. Conclusion

We have presented and analyzed a set of notorious teasers in Probability Theory. We pointed out some analogies among the different problems, and among their solutions. They all ask for the probability that the objects of a certain pair both have some property when information exists that at least one of them does. All, implicitly or explicitly, make reference to the manner in which that information was obtained, wherein lurks a potential source of difficulty.

We illustrated how different scenarios for obtaining some information yielded different solutions. In other words, the way we 'model' a problem is strongly dependent on the answer to the question: 'How was the information obtained?' On occasion, however, this dependence has been denied. Thus Neisser (1966), in a criticism of Freund's (1965) paper, said:

> 'The so-called 'puzzle' clearly states the information available. ... Hence, the result must be independent of the particular way in which the information was obtained ... Indeed, if the way in which information is obtained is allowed to influence the results, there is no end to possible modifications.' (p. 37).

Contrary to Neisser's claim, we showed that different ways of obtaining the selfsame information can significantly alter the revision of probability contingent upon it. All our problems highlighted the difference between knowing that at least one of two objects has some property on the basis of observing both, *versus* observing only one.

While it is sometimes natural or convenient to distinghuish between obtained information and the means whereby it was obtained, this distinction is often tenuous. Furthermore, the inclination to identify the conditioning event with the former automatically is fallacious. Such an identification can be justified only upon verification against the statistical experiment. Information cannot, as a rule, be divorced from its sources, and to do so can have devastating consequences.

The kind of problem in which the conditioning event does turn out to be identical to what is perceived as 'the information obtained' can only be found in textbooks. Consider a problem which asks for 'the probability of A *given* B'. This nonepistemic phrasing sidesteps the question of how the event B came to be known, since the term 'given' supplies the conditioning event, by definition. For example, the answer to the question: 'What is the probability that "Smith has two sons" *given* "Smith has at least one son"?' is (under the standard assumptions) unequivocally 1/3. Outside the never-never land of textbooks, however, conditioning events are not handed out on silver platters. They have to be inferred, determined, extracted. In other words, real-life problems (or textbook problems purporting to describe real life) need to be *modeled* before they can be solved formally. And for the selection of an appropriate model (i.e., probability space), the way in which information is obtained (i.e., the statistical experiment) is crucial.

We mentioned above that the sense of paradox which emanates from these problems is not completely resolved when we construct a successful model. Remember that in our presentation, the riddle posed by Problem 2 *versus* Problem 1 was: 'How can information which seems relevant (birth order) fail to affect the required probability?', whereas in Problem 4 *versus* Problem 3, it was: 'How can information which seems irrelevant (card suit) affect the required probability?' Once the correspondence between card suit and birth order becomes evident, a third riddle springs forth, namely: 'What's going on? *Is* birth order (card suit, name of lucky prisoner, card's upfacing color) relevant, or isn't it? Can it be one way for one problem, and another for a different, though almost identical, problem?' In fact, an effect which is startlingly similar to the reversible-figure effect in perception can be set up, whereby one can endlessly vacillate between two incompatible diagnoses of what the matter with these problems even is!

This lingering sense of paradox is due, we believe, to our intuitive notions of relevance. The puzzle which one perceives in these problems is largely determined by which of the following rules of thumb for relevance they highlight. (i) Information is relevant if it alters the conditioned sample space (as the addition of birth order narrowed it down from {BB, BG, GB} to {BB, BG}). (ii) Information is irrelevant if it doesn't matter which of its alternative values is supplied (as when stating the card suit as spades *versus* clubs).

Our problems have shown why neither of these criteria can serve as a definitive characterization of relevance. Sometimes a sample space may be altered upon the addition of information, but along with it the odds are altered as well, leaving the target probability unaffected. And sometimes the alternative possible values of the additional information are not really complementary events in the appropriate sample space; therefore, their symmetry need not lead to independence.

References

Beckenbach, E. F. (1970) Combinatorics for school mathematics curricula. In L. Råde (ed.), *Th. Teaching of Probability and Statistics.* Stockholm, Almqvist and Wiksell.

Betteley, I. G. (1979) Ambiguity in probability problems. *Math. Gaz., 63,* 16–19.

Falk, R. (1978) *Problems in Probability and Statistics,* Part II, (in Hebrew). Jerusalem, Akademon 68–69.

Frauenthal, J. C. and Saaty, T. L. (1979) Foresight–insight–hindsight. *Two-year Coll. Math. J., 10* 245–254.

Freund, J. E. (1965) Puzzle or paradox? *Amer. Stat., 19,* 29.

Gamow, G. and Stern, M. (1958) *Puzzle-math.* New York, Viking Press, 33–42.

Gardner, M. (1959) *The Scientific American Book of Mathematical Puzzles and Diversions.* New York, Simon and Schuster, 49–51.

Gridgeman, N. T. (1967) Letter. *Amer. Stat., 21,* 38.

Jeffrey, R. C. (1968) The whole truth. *Synth., 18,* 24–27.

Mosteller, F. (1965) *Fifty Challenging Problems in Probability with Solutions.* Reading, MA, Addison-Wesley.

Neisser, H. (1966) Letter. *Amer. Stat., 20,* 37.

Résumé

L'article porte sur un groupe de casse-tête bien connus en probabilité. Dans tous ces problèmes la question posée concerne la probabilité pour deux éléments d'une paire de posséder une caractéristique étant donné l'information qu'un des éléments de la paire possède. L'utilisation des probabilités conditionnelles permet de résoudre ces problèmes mais leur caractérisation correcte est difficile. En effet, ces problèmes soulignent l'importance de la manière dont l'information est obtenue. Une espèce de probabilité modélisant les problèmes verbaux devrait permettre la représentation du résultat obtenu *et* de l'expérience statistique qui le produit. On donne les raisons psychologiques qui rendent compte de la nature fallacieuse de ces problèmes et des moyens pratiques pour en venir à bout.