

GUILT BY STATISTICAL ASSOCIATION: REVISITING THE PROSECUTOR'S FALLACY AND THE INTERROGATOR'S FALLACY

The “prosecutor’s fallacy” and the “interrogator’s fallacy” are names for two different and very specific ways of inferring a suspect’s guilt. Although I do agree that these two kinds of reasoning involve some elementary logical mistakes, I think that in exposing these fallacies the critics sometimes introduce serious misconceptions of their own. They try to educate the wider public about how to deal with these probability issues but in the end they just manage to muddy the waters even more with their logically flawed analysis. In particular, I will try to show that in each of these two situations the critics fail to recognize some legitimate probabilistic indicators of guilt that remain available for a nonfallacious inference. My diagnosis of the problems in the two cases will be structurally very similar, and this explains why they are being discussed together in this article.

I. THE PROSECUTOR’S FALLACY

The prosecutor’s fallacy is a conflation of two conditional probabilities. Let M be the statement that a defendant’s DNA matches the DNA of the perpetrator (obtained from the crime scene), and let I be the statement that the defendant is innocent. If $p(M/I)$ is very low (that is, it is very improbable that there would be a DNA match if the defendant is innocent), this is sometimes taken to mean that $p(I/M)$ is also very low (that is, that it is very improbable that the defendant is innocent if there is a DNA match). This is an obvious mistake because it is clear that the two probabilities are logically quite distinct and that neither can be directly inferred from the other.

The confusion of $p(M/I)$ and $p(I/M)$ is called the “prosecutor’s fallacy” because the prosecutor is often interested in establishing that $p(I/M)$ is very low (which would indeed amount to strong evidence that the defendant is guilty), but occasionally what he instead presents to the court as incriminating evidence is just the fact that $p(M/I)$ is very low, which *in itself* does not point to the guilt of the defendant at all.

It is easy to expose this error but unfortunately it is also easy to commit another error in the process. Sometimes the criticism of the prosecutor’s fallacy is taken too far and the evidential relevance of the DNA match for the defendant’s guilt is not recognized even when it does have probative value. In some cases, surprisingly, even distinguished statisticians fall into this trap.

Let us start with Gerd Gigerenzer, one of the currently most prominent experts in the application of probability to practical problems. In explaining the perils of the prosecutor's fallacy, he discusses a criminal case from Germany in which two pieces of evidence were adduced against a defendant.¹ First, the defendant's blood matched the blood found under the fingernails of a murdered woman, with this particular blood type being shared by 17.3 percent of Germans. Second, a different blood type was found on the defendant's boots and it matched the blood of the woman, with this blood type being shared by 15.7 percent of Germans. Now how probable is it that the defendant is guilty if only this evidence is taken into account? Gigerenzer addresses the question and offers an answer that is inadequate on multiple levels.

He starts with estimating $p(M/I)$, that is, the probability of the blood match, given that the defendant is innocent. Gigerenzer suggests that we calculate the probability of the double match (first, between the suspect's blood and the blood under the woman's fingernails, and second, between the blood on the suspect's boots and the woman's blood) by simply multiplying the probabilities of each match (*ibid.*, p. 157). But this is a patently wrong approach to the second match. The fact that 15.7 percent of people share the same blood with the victim surely cannot mean that the probability of an innocent person having this blood type *on his boots* is also 15.7 percent! It must be much, much lower, simply because it is extremely rare that an innocent person has someone else's blood on his boots (or anywhere on his clothes, for that matter).

After allowing this ill-considered multiplication of probabilities (17.3 percent \times 15.7 percent) Gigerenzer concludes that the probability of the double match (given innocence) is 2.7 percent. So $p(M/I) = 0.027$. But what is the numerical value of $p(I/M)$? Those who commit the prosecutor's fallacy would say that it is the same, and from this it would follow that the probability of guilt, $p(\sim I/M)$, is overwhelming, namely, higher than 0.97.

This is of course wrong but even if we forget the problems with Gigerenzer's unhappy multiplication of probabilities his estimate of $p(\sim I/M)$ is still not correct. He first assumes that any of roughly 100,000 men in the city could have committed the crime. Of all these men, 2.7 percent will have the double blood match, which means around 2,700 of them. Now since the defendant belongs to a group

¹Gigerenzer, *Reckoning with Risk: Learning to Live with Uncertainty* (New York: Penguin, 2002), pp. 156–58.

of 2,700 men, and one member of that group is the perpetrator, the probability that the defendant is the murderer is 1 in 2,700. Therefore, according to him, $p(\sim I/M)$ is not 0.97 but 0.0004.

Interestingly, in another text² in which he discusses the same case, Gigerenzer is more careful and he realizes that the probability of someone's having the victim's blood on his boots cannot be equal to the proportion of people with that particular blood type in general population. He is aware that this probability must be much lower (because not everyone has someone else's blood on him) but, oddly enough, he now widens the range of suspects from the previous 100,000 to 10 million, and by fudging the numbers in this way he makes sure that the probability of the suspect's guilt given the DNA match again comes out very low.

Before I point out the central mistake in Gigerenzer's reasoning let me present a similar analysis of another case (*R. v Adams*) that is offered by the Oxford statistician Peter Donnelly.³ A man was convicted of raping a woman, and again the only incriminating evidence against him was the fact that his DNA profile matched the DNA evidence from the crime scene. Donnelly claims that even if this DNA type were so rare that it is found only in one person in 20 million, this would still not make it more probable that the defendant was guilty than that he was not.

His argument is essentially the same as Gigerenzer's. From the assumption that $p(M/I)$ is 1 in 20 million Donnelly concludes: "If you believe that number, then on average there will be 2 or 3 people in Britain whose DNA it could be, and probably no more than 6 or 7" (*ibid.*, p. 48). It would seem then that despite the extremely low probability of the DNA match (on the assumption of innocence), the probability of the defendant's guilt is still not very high even after it is established that his DNA matches the DNA of the perpetrator.

Although Donnelly does not provide a mathematical derivation of that result, there is a simple Bayesian formula that underlies his argument. The goal is to obtain the probability of guilt, or $p(G/M)$, given the DNA match. The calculation is straightforward because all four necessary elements are known.

$p(G)$ = the probability of an arbitrarily chosen person being guilty = 1 in 60 million (all people in Britain)

$p(\sim G)$ = the probability of an arbitrarily chosen person not being guilty = $1 - p(G)$, that is, very close to 1

²Gigerenzer, "Ecological Intelligence: An Adaptation for Frequencies," in D.D. Cummins and C. Allen, eds., *The Evolution of Mind* (New York: Oxford, 1998), pp. 9–29.

³Donnelly, "Appealing Statistics," *Significance*, 11 (2005): 46–48.

$p(M/G)$ = the probability of the DNA match if the person is guilty = 1

$p(M/\sim G)$ = the probability of the DNA match if the person is not guilty
= 1 in 20 million

$$(1) \quad p(G/M) = \frac{p(G) \times p(M/G)}{p(G) \times p(M/G) + p(\sim G) \times p(M/\sim G)}$$

$$= \frac{\frac{1}{60m} \times 1}{\frac{1}{60m} \times 1 + \approx 1 \times \frac{1}{20m}} = 0.25$$

Donnelly is aware, of course, that he neglects some relevant considerations, and that many people in Britain would obviously be excluded from suspicion on a number of grounds (like gender or age). Yet although he explicitly warns that the class of potential suspects is for this reason much narrower than the whole population of the United Kingdom and although it is quite easy to introduce this correction, Donnelly inexplicably proceeds to work out his calculation on the basis of this very unrealistic assumption. The result (the probability of 0.25, obtained in (1)), is a serious underestimation of the true probability of guilt. For if we first exclude all women (30 million) and then also a number of men that are ruled out on grounds like age and so on (probably around 20 million), we are left with the group of potential suspects of approximately 10 million. As equation (2) shows, doing the calculation on the basis of this much more plausible assumption, $p(G/M)$ turns out to be considerably higher than previously, namely 0.66.

$$(2) \quad p(G/M) = \frac{p(G) \times p(M/G)}{p(G) \times p(M/G) + p(\sim G) \times p(M/\sim G)}$$

$$= \frac{\frac{1}{10m} \times 1}{\frac{1}{10m} \times 1 + \approx 1 \times \frac{1}{20m}} = 0.66$$

Yet there is a more serious mistake lurking in Donnelly's and Gigerenzer's approach to the prosecutor's fallacy. Although they are right that the simple transposition of two probabilities, $p(M/I)$ and $p(I/M)$, is a fallacy, I will argue that in each of their two cases the DNA match actually does constitute much stronger evidence that the defendant is guilty than they recognize. The basic mistake in their analysis consists in disregarding a crucial aspect of the situation, which should have been taken into account and which indeed significantly lowers the estimated value of $p(I/M)$, and consequently increases the probability of the defendant's guilt.

The fact that in *R. v Adams* the defendant's DNA was found to match the DNA from the crime scene entails that his DNA *was known to the police*. That is, it was available in the database that the police used in looking for a match. And precisely this piece of information is the dog that barked loudly but was ignored by both Gigerenzer and Donnelly. If it were true that the police databases contain DNA specimens of different people based on completely random sampling, the reasoning of the two statisticians would be perfectly fine. But clearly this is not how things are. It is common knowledge that the DNA included in police files predominantly belongs to individuals with previous convictions or arrests. And, most importantly, it is well established that the probability of these people committing a crime is significantly higher than for the general population. For instance, in the research on this matter that was recently conducted for the British Home Office, it is stated as completely uncontroversial that "generally the best predictor of future offending is a previous history of offending."⁴

So the situation is more complicated than presented in (1) and (2). For estimating the probability of the defendant's guilt *all things considered* there are *two* facts that have to be taken into consideration, not just one. Besides *M* (the fact of the DNA match) there is also *D* (the fact that the defendant's DNA was in the police database). In other words, we are not looking for $p(G/M)$ but actually for $p(G/M \& D)$. Here is an appropriate equation for that purpose.

(3)

$$p(G/M \& D) = \frac{p(G/D) \times p(M/G \& D)}{p(G/D) \times p(M/G \& D) + p(\sim G/D) \times p(M/\sim G \& D)}$$

There are four probabilities on the right side of the equation. Using Donnelly's example again, the magnitudes of two of these probabilities are already known:

$p(M/G \& D)$ = the probability of DNA match given that one is guilty and in database = 1

$p(M/\sim G \& D)$ = the probability of DNA match given that one is innocent and in database = 1 in 20 million

The probability of *M* given *G* and *D* is 1 because it was assumed from the beginning that the DNA collected at the crime scene did belong to the perpetrator. And the probability of *M* given $\sim G$ and *D* is 1 in 20 million because if one is not guilty then the probability

⁴ B. Powis, *Offenders' Risk of Serious Harm: A Literature Review* (London, UK: Home Office, 2002), p. v.

of one's DNA matching the perpetrator's DNA must be the same, *whether one's DNA is in a police database or not.*

How about the two remaining probabilities: $p(G/D)$ and $p(\sim G/D)$? They have to add up to one, so we have to assign the value only to one of them. For the sake of the discussion, let us suppose first that the probability that a person with previous conviction(s) has 100 times higher probability of committing a crime than a randomly chosen person from the general population.⁵ Then since the probability of an arbitrarily picked out individual committing this particular crime was assumed to be 1 in 10 million, we obtain the following values for the remaining two unknowns:

$$p(G/D) = \text{the probability of guilt given that one's DNA is in police database} = 1 \text{ in } 100,000$$

$$p(\sim G/D) = \text{the probability of innocence given that one's DNA is in database} = 99,999 \text{ in } 100,000$$

Substituting these values into equation (3) yields 0.99 as the value of $p(G/M\&D)$:

$$(4) \quad p(G/M \& D) = \frac{\frac{1}{100,000} \times 1}{\frac{1}{100,000} \times 1 + \frac{99,999}{100,000} \times \frac{1}{20,000,000}} = 0.99$$

Alternatively, if we suppose that $p(G/D)$ is only fifty times (rather than a hundred times) higher than $p(G)$, equation (4) would still give the same answer for $p(G/M\&D)$ as before: 0.99. Even if $p(G/D)$ is set to be as low as just ten times higher than $p(G)$, $p(G/M\&D)$ would continue to be higher than 0.95.

In Gigerenzer's case as well, an analogical correction produces a similar surge in the key probability. Recall that in his example the value calculated for $p(G/M)$ was 0.0004. After necessary amendments are made in order to correct for the ludicrous multiplication of probabilities and after it is taken into account (in a similar way as with Donnelly's example) that those whose data are in police databases are significantly more likely to commit a crime, again the estimate of $p(G/M\&D)$ jumps from 0.0004 to values higher than 0.9.

⁵ If some readers feel that this estimate might be a wild and unrealistic exaggeration let me point out that, for example, according to criminal statistics from England and Wales, a person convicted of a sexual offense (the crime that is relevant for our context) is between 80 and 250 times more likely to be convicted of a similar act in the first two years after being released than some randomly chosen individual. The source: M. Redmayne, "The Relevance of Bad Character," *Cambridge Law Journal*, LXI (2002): 695.

The fallacious reasoning that is criticized here has even crept into the pages of *Nature*. Statisticians David Balding and Donnelly consider a hypothetical case in which $p(M/I)$ is 0.000001, and they try to show that under certain assumptions $p(I/M)$ would still be relatively high, or 0.33 to be exact. Here is their reasoning: "There may be many individuals who, *if not for the DNA evidence, would be as likely to be the culprit as the defendant*. If there were 500,000 such individuals, for example in a large city, the probability that the defendant is innocent after taking the DNA evidence into account would be at least one-third."⁶

The italicized phrase is wrong because it implies that in the absence of the DNA match the defendant would be as likely to be the culprit as any member of the general population that possesses the general characteristics ascribed to the culprit (for example, male, age between 20 and 40, and so on). But in fact, as explained above, even without the DNA match the defendant's probability of guilt is already significantly higher than for the average person with those given characteristics, due to the presence of his DNA in police databases.⁷

Someone might object here that the increased probability of committing a crime that is inferred from the presence of one's DNA in police databases (and ultimately from the fact that these available DNA specimens massively belong to previous offenders) would be inadmissible evidence in a court of law and that Gigerenzer and Donnelly therefore justifiably ignored this information. There are three reasons why this defense is not convincing.

First, the previous offender status is not universally inadmissible for determining guilt. For instance, in some U.S. states (for example, Wisconsin) the prosecutor is allowed to introduce this kind of evidence in connection with some kinds of sexual offenses. And remember that both cases discussed by the statisticians involve rape. Moreover, some scholars argue that the criminal justice is already rapidly moving in the direction of allowing the evidence of bad character in criminal trials: "The winds of change look to be blowing towards radical reform of the policy which excludes evidence of a defendant's bad character. In the near future, evidence of previous misconduct may be presented to juries much more often than at present."⁸

⁶ Balding and Donnelly, "How Convincing Is DNA Evidence?" *Nature*, CCCLXVIII (1994): 285 (italics added). Cf. Balding and Donnelly, "The Prosecutor's Fallacy and DNA Evidence," *Criminal Law Review* (1994): 719.

⁷ Indeed, in the aforementioned rape case (*R. v Adams*), in which Donnelly was a witness for the defense, the defendant's DNA also became available to the police only after he was convicted of another, separate sexual assault. See "Rapist Jailed on Sole Evidence of DNA Database," *The Times* (September 14, 1996).

⁸ Redmayne, "The Relevance of Bad Character," *Cambridge Law Journal*, LXI (2002): 714.

Second, the fact that this kind of evidence is formally inadmissible to the court does not mean that the jurors will not use it tacitly in their deliberations. After all, in reaching a decision they are asked to use their common sense, and it is part of common-sense knowledge that DNA databases largely contain samples of previous offenders with increased crime risk. Therefore, any statistician educating the public about the prosecutor's fallacy should address the impact of that piece of empirical knowledge on the estimation of relevant probabilities (even if he decides in the end to advise the reader to suppress this information in the *legal* determination of guilt). If this issue is completely ignored, the probability analysis offered by the statisticians will be unconvincing since it really omits a crucial factor that is highly relevant for the correct calculation of probabilities. In other words, it may well be that the statisticians' solution encounters resistance and looks so counterintuitive and unacceptable not because people are confused about probabilities but rather because this "solution" proposed by the statisticians is just *wrong*, all things considered.

Third, someone's guilt is often an issue in an out-of-court context. The correct handling of probabilities is very important there as well, and it is not restricted by usual legal limitations. Suppose that you learn that your children's baby-sitter has been interrogated by the police in a child-molesting case because his DNA matched the DNA collected at the crime scene. Would it not be unwise and perhaps even irresponsible in this situation to decide whether to keep this baby-sitter or not by completely ignoring the fact that the police usually run DNA searches through its DNA database that contains samples from a disproportionately high number of potentially dangerous individuals?

II. THE INTERROGATOR'S FALLACY

The name "interrogator's fallacy" was introduced into the literature by Robert A. J. Matthews.⁹ He defines this fallacy as "the error of assuming that confessional evidence can never reduce the probability of guilt." Now it is true that we think that the existence of confession normally increases the probability of guilt. And we are right about this. But Matthews is also right that under certain circumstances the probabilistic impact of a confession could actually go in the opposite direction and *make the guilt less likely*.

The crucial question is when exactly the fact of a confession is an indicator of guilt and when not. Matthews gives the following answer:

⁹ Matthews, "Inference with Legal Evidence: Common Sense Is Necessary but Not Sufficient," *Medicine, Science and Law*, XLIV (2004): 189–92; Matthews, "The Interrogator's Fallacy," *Bulletin of the Institute of Mathematics and its Applications*, xxxi (1994): 3–5.

confession increases the probability of guilt *only* when the probability of confessing given that the person is guilty is higher than the probability of confessing given that the person is innocent. If G refers to the statement that a person is guilty and C to the statement that the person confesses, then what Matthews is saying is that $p(G/C) > p(G/\sim C)$ only under the condition that $p(C/G) > p(C/\sim G)$.

The mathematician Ian Stewart defended the same claim in an article "The Interrogator's Fallacy" in *Scientific American* (in the widely read column "Mathematical Recreations" that he inherited from Martin Gardner). Stewart also argues that "the existence of a confession increases the probability of guilt if and only if an innocent person is less likely to confess than a guilty one."¹⁰ Matthews and Stewart have no doubts about the truth of their claim because they think that it is derivable from some elementary probability formulas. And indeed, they can rely on the following statement that is known as the odds form of the Bayes's theorem:

$$(5) \quad \frac{p(G/C)}{p(\sim G/C)} = \frac{p(G)}{p(\sim G)} \times \frac{p(C/G)}{p(C/\sim G)}$$

On the left side of (5) is $p(G/C)/p(\sim G/C)$, the ratio of so-called *posterior* probabilities of guilt and innocence (that is, probabilities *after* the confession is taken into account). It is equal to the product of two other ratios on the right side of (5). The ratio $p(G)/p(\sim G)$ represents the ratio of *prior* probabilities of guilt and innocence (that is, probabilities of guilt and innocence *before* the fact of confession is factored in). The last ratio, $p(C/G)/p(C/\sim G)$, is called the likelihood ratio and obviously it will determine what impact the new evidence (confession) will have on the probabilities of guilt and innocence. In particular, the posterior probability of guilt will be higher than the prior probability of guilt only if the likelihood ratio is larger than 1. But this is exactly what Matthews and Stewart are telling us.

In addition, they argue that in some situations the likelihood ratio, $p(C/G)/p(C/\sim G)$, will be smaller than 1, and that in these cases a confession will actually become an indicator of *innocence*. Their example: due to their training and fanaticism, hardened terrorists are in a better position to resist rough interrogation than innocent people who, being much more compliant and suggestible, are unable to cope with the

¹⁰ Stewart, "The Interrogator's Fallacy," *Scientific American* (September 1996): 134–36; Stewart, *The Magical Maze: Seeing the World through Mathematical Eyes* (New York: Wiley, 1997), pp. 90–93.

unpleasant situation and often end up by confessing to a crime that they did not commit.

I agree, of course, that (5) is a logical consequence of the axioms of probability theory. I also agree that in the example with terrorists the likelihood ratio could indeed be lower than 1. But I do not agree that it necessarily follows from all this that the existence of confession in that situation would be evidence against guilt. Contrary to what Matthews and Stewart are saying, I think that *even if an innocent person is more likely to confess than a guilty one, the confession could still raise the probability of guilt.*

How can this be? Such an idea seems to be a logical impossibility if we look at (5). But my point is that equation (5) may not be the right formula to use here. There could be another empirical fact that is relevant for the probability of guilt and that escaped our attention. In that case we would obviously have to use a more complicated formula that would incorporate that fact.

Recall how in the case of the prosecutor's fallacy we initially focused exclusively on the DNA match and how at first we completely missed the fact that the defendant's DNA match presupposes the existence of the defendant's DNA in police databases, a very relevant thing for estimating the probability of guilt.

Well, I think that Matthews and Stewart are again missing the possible (and moreover probable) existence of something very important in this context: the very fact of interrogation. Why would this be relevant? Simply because there is an obvious possibility that among those interrogated by the police there is a higher proportion of criminals than in the general population.

Since we are examining here the evidential impact of confession *that is the result of interrogation*, the probability we are looking for is not really the probability of X being guilty given that he confessed, but rather the probability of X being guilty given that he confessed *and that he was interrogated*. Therefore, the ratio of posterior probabilities should be expressed in the following more comprehensive way (with ' I ' standing for "having been interrogated"):

$$(6) \quad \frac{p(G/C \ \& \ I)}{p(\sim G/C \ \& \ I)} = \frac{p(G/I)}{p(\sim G/I)} \times \frac{p(C/G \ \& \ I)}{p(C/\sim G \ \& \ I)}$$

... or even better, by expanding further the first probability ratio on the right side of (6):

$$(7) \quad \frac{p(G/C \ \& \ I)}{p(\sim G/C \ \& \ I)} = \frac{p(G)}{p(\sim G)} \times \frac{p(I/G)}{p(I/\sim G)} \times \frac{p(C/G \ \& \ I)}{p(C/\sim G \ \& \ I)}$$

To see why the first ratio on the right side of (6) is equivalent to the product of the first two ratios on the right side of (7), consider first the following two elementary truths in probability:

$$(8) \quad p(G/I) = \frac{p(G) \times p(I/G)}{p(I)}$$

$$(9) \quad p(\sim G/I) = \frac{p(\sim G) \times p(I/\sim G)}{p(I)}$$

Now if (8) is divided by (9), the magnitude $p(I)$ in the denominators on the right side cancels out and we obtain:

$$(10) \quad \frac{p(G/I)}{p(\sim G/I)} = \frac{p(G)}{p(\sim G)} \times \frac{p(I/G)}{p(I/\sim G)}$$

Equation (10) proves that (6) and (7) are indeed equivalent.

Let us compare equation (5), used by Matthews and Stewart, with equation (7) that gives a more complete picture of the situation. On the left side in each case is the ratio of *posterior* probabilities of guilt and innocence, with the difference that in (5) these probabilities are conditioned only on confession, while in (7) they are conditioned on both confession and interrogation. The first ratio on the right side is the same in both (5) and (7): the ratio of prior probabilities of G and $\sim G$. The last ratio on the right side is also similar in the two cases: the probability of confession given guilt (and interrogation) divided by the probability of confession given innocence (and interrogation).

There is a completely new element in (7), however, and it is the middle ratio on the right side: $p(I/G)/p(I/\sim G)$. In words, this is the probability of being interrogated by police given that the person is guilty divided by the probability of being interrogated by police given that the person is innocent. We have good reasons to believe that this ratio is higher than 1, unless the police are very incompetent. For it stands to reason to expect that the police will more probably select for interrogation someone who is guilty than someone who is innocent. Furthermore, it seems reasonable to assume that the rougher the interrogation methods used by the police, the more important it will be for them to make sure that those selected for that treatment are indeed guilty.

It is precisely this new element in (7) that reveals what is wrong with Matthews's and Stewart's claim that the probability of guilt can increase after confession *only if* $p(C/G)$ is higher than $p(C/\sim G)$.

Here is a counterexample to their claim. Let us accept their assumption that hardened criminals less frequently confess under police interrogation than innocent people. For the sake of concreteness, let us

suppose that, say, only 40 percent of guilty people confess while 60 percent of innocents do. Since this makes the probability of confession given guilt *lower* than the probability of confession given innocence, it would seem to follow (as we are told by Matthews and Stewart) that now the existence of a confession cannot increase the probability of guilt from what it was before the confession. But suppose further that the police are 9 times more likely to interrogate a guilty person than an innocent individual. Finally, suppose that the prior probability of a particular person (let us call him John) being guilty, based on all the evidence before the confession, is 0.75.

Here are the probabilities that apply to the situation as described:

$$p(C/G \& I) = 0.4$$

$$p(C/\sim G \& I) = 0.6$$

$$p(G) = 0.75$$

$$p(\sim G) = 0.25$$

$$p(I/G)/p(I/\sim G) = 9$$

Substituting these values into equation (7) yields:

$$(11) \quad \frac{p(G/C \& I)}{p(\sim G/C \& I)} = \frac{p(G)}{p(\sim G)} \times \frac{p(I/G)}{p(I/\sim G)} \times \frac{p(C/G \& I)}{p(C/\sim G \& I)}$$

$$= \frac{0.75}{0.25} \times 9 \times \frac{0.4}{0.6} = 3 \times 9 \times \frac{2}{3} = 18$$

Equation (11) tells us that the probability of John being guilty, *all things considered*, is 18 times higher than the probability that he is innocent. Since he must be either guilty or innocent, we can use this ratio to obtain the value for the posterior probability of his guilt. The probability that he is guilty given that he confessed (and that he was interrogated) is 18/19, or 0.95.

Before confession, the probability of John's guilt was 0.75. After confession, the probability of his guilt increases to 0.95 *despite the fact that the probability that he confesses if he is guilty is lower than the probability that he confesses if he is innocent*. This completes the proof that Matthews and Stewart are wrong when they insist that confession can be evidence of guilt *only if* the probability of confession given guilt is higher than the probability of confession given innocence.

III. CONCLUSION

We have seen that some experts in mathematical statistics have given seriously misguided advice about how to avoid fallacious reasoning about probabilities. They try to cure us from the prosecutor's fallacy

and the interrogator's fallacy but it turns out that their medicine is sometimes worse than the disease.

Yet as grave as their lapses of judgment may have been, I surely cannot on that ground alone endorse W.H. Auden's stern admonition "Thou shalt not sit with statisticians..." But a milder version of that warning might be in order. If you really must sit with statisticians, do not let them always persuade you too quickly and do not give up your common sense opinion as soon as it clashes with their analysis. For as demonstrated in this article, on some rare occasions your gut intuition about probabilities may happen to be more trustworthy than what these professors of statistics are trying to teach you.

CARMEN DE MACEDO

Hong Kong