Parabola Volume 28, Issue 1 (1992)

CURRENCY, CROWS AND UNEXPECTED EXAMINATIONS

David Angell*

Part 1: problems.

You have (let's imagine) five maths lessons per week, one each day. One week, at the end of Friday's class, the teacher announces, "In one of next week's lessons there will be an unexpected examination. By saying that it will be unexpected, I mean that you will not be able to deduce when it will be until the day it actually happens..."

... on Sunday evening you are to be found leaning upon your desk, repeating over and over the ancient student incantation, "Why didn't I study last week when I had the chance?" There might be some hope of passing if only the exam turns out to be on Friday...

Suddenly a thought strikes. *Could* the exam be on Friday? If so, then after Thursday's class I wouldn't have had the test; so I would *know* it was to be on Friday. But this is impossible, because the teacher said that I wouldn't be able to deduce the date of the exam in advance. So it looks as if the exam can't be on Friday! "Great," you mutter to yourself, "I can't prove the Binomial Theorem, but I can prove I'm going to totally fail this test. I wonder if I get any marks for doing that?"

After a few minutes' quiet desperation, "What about Thursday?" you ask yourself. "Well, if the test were to be held on Thursday, then at the end of Wednesday's class I would be certain that either Thursday or Friday was to be the day; but I already know that Friday is impossible; so it must be Thursday! Once again this contradicts the teacher's statement that I wouldn't know the day of the exam until it actually arrives, and therefore proves that Thursday is impossible. The same argument disposes of Wednesday and Tuesday. So the test must be on Monday – but since I have already deduced it on Sunday, this too is impossible!" Things are looking up.

That night, having decided that it's not worth while studying for a non-existent test, you dream of sitting in your maths class on Monday morning, hand raised high. "You have

* David is a pure mathematician at the University of New South Wales.

a question?" asks the teacher. Carefully you explain why the unexpected examination cannot be given at all. "Oh dear," says the teacher, "I hadn't thought of that," and promptly vanishes in a puff of logic...

In Albury-Wodonga, on the New South Wales-Victoria border, they have too many problems paying bills to spend time worrying about maths tests. In an attempt to solve their economic difficulties the states of NSW and Victoria banned the use of the Australian dollar, and each introduced their own currency, known as the NSWollar and the Vicdollar. This would not be much of a problem except that the two states started to develop serious disagreements between themselves. Eventually the NSW government declared that in NSW, a Vicdollar would be worth only 90 NSW cents; and Victoria naturally retaliated by proclaiming that in Victoria, one NSWollar would be worth only 90 Victorian cents.

But as the saying goes, every problem merely conceals an opportunity. A resident of Albury, hearing the advice from Canberra that Australia should "spend its way out of the recession", hit upon a very good opportunity indeed to spend his way out of his own personal recession. He took a ten NSWollar note, went to an Albury bakery and bought a loaf of bread worth one NSWollar. He paid for the bread with the ten NSWollar note, and received back a ten Vicdollar note; this being the correct change since it was worth nine NSWollars. He then crossed the border to a Wodonga milk bar and bought a one Vicdollar carton of milk; paid for it with the ten Vicdollar note; and was given (correctly) a ten NSWollar note as change. On returning home, he had all of his original money, and had therefore done his shopping for free! The baker was one NSWollar richer, having ten NSWollars instead of ten Vicdollars; and likewise, the milk bar owner was one Vicdollar richer. So who paid for the bread and milk?

Those who get headaches from trying to outguess their maths teacher, or from juggling NSWollars and Vicdollars, might instead take up ornithology, a subject totally free of annoying logical problems. Or is it?

After years of thought an ornithologist has come up with a startling theory: "all crows are black". To check the accuracy of this theory is a simple matter of observation. The ornithologist takes a trip out to the bush, or anywhere else crows are to be found, identifies certain objects as crows, and confirms that their colour is indeed black. Every such observation supports her theory, and by now she has collected a good many records. Of course, if by any chance she were to identify an object as a crow and then find that its colour was not black, the whole theory would fall to the ground. So far, however, this hasn't happened.

One day the weather is cold and wet, and the scientist, feeling disinclined to visit the bush in such conditions, sits at home wondering how best to use the time. There are no crows in her house, so direct research is out of the question... unless... It dawns upon her that the problem could be approached differently. The statement "all crows are black" is, from a logical point of view, exactly the same as "all non- black things are not crows". So instead of identifying crows and examining them for blackness, it seems just as valid to identify non-black things and examine them for non-crowiness. Glancing around her she finds a rich source of evidence: on a nearby table stands a bowl of fruit. In it there sits an object which is not black, but yellow; on examination it is found to be not a crow, but a banana. This is an observation in favour of the black crow hypothesis. A quick glance reveals another non-black object which, on the evidence of a taste test, appears to be a strawberry rather than a crow. (Just as well.) Marvellous! – still more support for the theory!

Like all good scientists, our ornithologist friend takes some time out to consider her methodology. It appears ridiculous to say that a banana can have any relevance to ornithological research. Yet the logic seems watertight. Observation of the banana clearly supports the hypothesis "all non-black things are not crows", and therefore must also support the proposal "all crows are black", as this is merely a different way of saying exactly the same thing.

Her meditations are interrupted by a violent knocking at the door. She opens the door to see a fellow ornithologist, his coat streaming with rain. With immense excitement he exclaims, "After years of thought I've come up with a startling theory: 'all crows are white'! And," he continues, advancing towards the fruit bowl, "I know *exactly* where to look for evidence..."

You are sitting in your maths class on Monday morning, hand raised high. "You have a question?" asks the teacher. "Well I'm afraid it will have to wait until after the test..."

The test is not a highlight of your mathematical career. Afterwards, you reconsider the situation. The test was given, despite your "proof" that it could not possibly be. Furthermore, the teacher's claim that it would be unexpected turned out to be clearly and horribly true. What went wrong?

Part 2: answers, sort of.

One way to deal with the currency paradox is to note that (for example) 9 NSWollars are worth 10 Vicdollars in NSW, but only 8.10 Vicdollars in Victoria. So we could simply shrug our shoulders and say that conditions in the two states are inconsistent, and that it is meaningless to speak of trade between them. This is a rather unsatisfactory solution, however, since the proposed regulation of currency values is exactly the sort of thing we could imagine being done by two opposing governments. It seems pointless simply to assert that it couldn't happen.

A more significant fact that we might observe is that apart from getting his bread and milk, the shopper has done one other thing: he has moved a ten Vicdollar note from NSW back to Victoria, and a ten NSWollar note back home to NSW. Now, nobody has any incentive to move notes in the reverse direction since they lose 10% of their value, so if this sort of process is carried out often enough, there will eventually be no more Victorian money in NSW, or no more NSW money in Victoria. Free trans-border shopping will then be impossible. Students of economics can probably think of other factors (for example, variation in prices of commodities in the two towns) which would after some time cause all the paradoxes to disappear. These arguments, however, are not much more satisfactory than the last. They prove that the current situation in Albury-Wodonga cannot continue for ever, but they do nothing to explain what is happening at present.

The idea of moving money, however, may suggest a better solution to the problem. How did the ten Vicdollar note get to NSW in the first place? It must have been brought from Victoria (possibly a long time ago), and the person who brought it would have lost money; for the note would have purchased, say, ten cartons of milk in Victoria, but only nine in NSW. This person, in effect, paid for the milk, and similarly, some other person paid in advance for the bread by moving a NSWollar to Victoria.

One problem (at least!) remains. What if the people carrying money interstate in the previous paragraph did so *before* the laws came into force fixing the values of NSWollars and Vicdollars in each state? Then these people did not lose any money, since the two currencies were of equal value. So who in this situation would have paid for the bread and milk? I'll leave this one for you to think about yourselves.

The problem in the ornithological tale is to explain the discrepancy between the apparently clear logic on the one hand, and on the other our commonsense conception of "supporting evidence", which tells us that a yellow banana has no possible relevance whatever to a question about the colour of crows.

Let's deal with the final episode first, as it is perhaps not too difficult to resolve. The suggestion made is that the same observation (the bowl of fruit) can be used as evidence for two contradictory conclusions: "all crows are white" and "all crows are black". However, the introduction of white crows is merely (so to speak) a red herring. The yellow banana more specifically supports the statement "all yellow things are not crows", which is logically the same as "all crows are non- yellow". It is clear that if we ever gained enough evidence to be confident that this last statement was true, it would still be quite possible that crows were all black, or all white. The hypotheses "all crows are black" and "all crows are white" are correctly said to be *contrary* (they cannot both be true, but could both be false) rather than *contradictory* (they can be neither both true nor both false). They *sound* contradictory since we are used to thinking of black and white as fundamental opposites. But in fact the present situation is simply a case where a set of observations is consistent with two different explanations: the two hypotheses concerning colour of crows may be advanced as "explanations" for the observation of a yellow non-crow. This particular observation, however, is not sufficiently specific to discriminate between the rival theories. A comparable situation occurred in Ancient Greek astronomy, where the available celestial data was insufficient to differentiate the theories "the sun moves round the earth" and "the earth moves round the sun". These two statements again are contrary but not contradictory. On the other hand, it is impossible to find an observation which will be evidence for each of the (truly contradictory) statements "all crows are black" and "not all crows are black".

To get some firmer ideas on the main problem, let's look at a smaller-scale example. I take a pack of playing cards, select some of them, and ask you, "All the red cards in my hand are diamonds: true or false?" You are allowed to demand that I show you what cards I hold of any one type (red, black, spades, hearts, diamonds or clubs), or of the negation of any one type (non-red, non-black, non-spades, and so on). I will then select all the appropriate cards, shuffle them (so that I can't try to fool you by placing them in a particular order), and show them to you one by one.

The obvious category for you to request is red cards. If you then see a series of diamonds turned up, you will become more and more confident that the correct answer is "true"; on the other hand, if you ever see a heart you will know that it is "false". However, like our first ornithologist, you could approach the problem differently by asking for all non-diamonds. If the cards I show you ever include a red one, then it will be a heart and you will answer "false". But if black after black turns up, it becomes more and more likely that there are no red cards to be shown among the non-diamonds, and therefore that all the red cards in my original hand must have been diamonds. It seems (I think) obvious that both ways are, in this example, equally valid.

Is this still true in the ornithologist's problem, where we have to consider the absolutely enormous number of non-black things in the world? Clearly a single observation will not give us much evidence; the question is whether or not it gives us any evidence at all. This is a harder problem, but I'm prepared to go out on a limb and suggest that the anti-commonsense solution is in fact the correct one: there is no paradox here, but merely a (well-justified!) feeling of surprise. A yellow banana truly is evidence in favour of the hypothesis "all crows are black" (and also for the hypothesis "all crows are white" - this is not contradictory, as I explained above), though the amount of evidence gained is so tiny as to be practically worthless. What do you think?

Finally we return to your classroom problems. This paradox was first noted in the

1940s, and attracted many attempted solutions in philosophical publications in the 1950s and early 60s. (By 1972 it had even reached year 8 schoolboys in Melbourne.) One suggested resolution depends on the word "deduce" in the teacher's statement. Suppose that what the teacher said is true. Then your reasoning above leads to the conclusion that it cannot be true after all, and so must be false. That is, there will not be an unexpected examination. This opens up two possibilities: either there will be no examination at all, or there will be one but it will not be unexpected. Since either of these cases could occur (the second occurs if the test turns up on Friday), it is strictly speaking impossible for you to *deduce* whether or not there will be an examination; and therefore, whenever it occurs (if it does) it will be unexpected in the sense in which the teacher defined that term. You may, of course, consider that you have strong grounds for *believing* the test will be on a certain date, but this is far from *deducing* that fact.

I find this explanation highly unconvincing. It fails to explain the fact that the teacher's statement turns out true when you have already (it seems) proved it false using a perfectly ordinary proof by contradiction. After all, the argument of the preceding paragraph scarcely even bothers to dispute your proof that an unexpected examination cannot be given. All it does is (perhaps) to shift the element of surprise away from the fact that the test took place and was truly unexpected, onto the fact that you were apparently able to prove this could not happen. I prefer the following explanation, though when you've read it you may feel that it is more of a cop-out than a real explanation.

Let's imagine that the teacher, instead of what was given in the original story, had said something a little different.

"(1) In one of next week's lessons there will be an exam.

(2) It will be unexpected; by which I mean that from statement (1) alone, you will not be able to deduce when the exam will be until the day it actually happens."

Now this situation involves no paradox, and the exam could in fact occur on any day except Friday. We can prove as before that Friday is not possible (check the details for yourself); however the argument will take us no further. For example, after Wednesday's class, if you have not had the exam already, you can deduce from statement (1) that it must be on Thursday or Friday. But a subsequent deduction that it is not on Friday (and so must be on Thursday) would depend on statement (2) as well as (1), and would therefore be inadmissible here. Certainly from the complete statement made by the teacher you will know that the test is to be on Thursday, and so it will hardly be "unexpected" in the ordinary sense of the word; but it *will* be unexpected according to the teacher's definition of this term, because you are not able to deduce that it will be on Thursday by use of statement (1) alone.

Our difficulties arise because the teacher's statement was in fact not the above, but the following:

"(1) In one of next week's lessons there will be an exam.

(2) It will be unexpected; by which I mean that from these two statements you will not be able to deduce when the exam will be until the day it actually happens."

I believe that this analysis shows that the paradox springs from a familiar source, *self-reference*. The new statement (2) asks us to deduce things about it, using the statement itself as a (partial) basis for our deductions, and thereby leads us into a logical whirlpool in which we eternally go round and round and get nowhere. It is well known that self-reference often (though not always) leads us into logical contradictions. The simplest example of this is the sentence, "This sentence is false." If it is false then it must be true; but if it is true it has to be false. It is generally agreed that statements of this sort – they can be expressed more precisely – are genuine paradoxes, which cannot be explained away but simply have to be avoided. I think that perhaps the unexpected examination falls into the same category. Rather a disappointing end to the tale – can you think of a better one?

Further reading.

Patrick Hughes and George Brecht. Vicious circles and infinity: an anthology of paradoxes. Penguin, 1989.

Martin Gardner. Mathematical puzzles and diversions. Penguin, 1991

Martin Gardner. The unexpected hanging and other mathematical diversions. University of Chicago Press, 1991.

Eugene P. Northrop. Riddles in mathematics. Pelican, 1974.