Solution

Week 2 (9/23/02)

Green-eyed dragons

Let's start with a smaller number of dragons, N, instead of one hundred, to get a feel for the problem.

If $N = 1$, and you tell this dragon that at least one of the dragons has green eyes, then you are simply telling him that he has green eyes, so he must turn into a sparrow at midnight.

If $N = 2$, let the dragons be called A and B. After your announcement that at least one of them has green eyes, A will think to himself, "If I do *not* have green eyes, then B can see that I don't, so B will conclude that she must have green eyes. She will therefore turn into a sparrow on the first midnight." Therefore, if B does not turn into a sparrow on the first midnight, then on the following day A will conclude that he himself must have green eyes, and so he will turn into a sparrow on the second midnight. The same thought process will occur for B, so they will both turn into sparrows on the second midnight.

If $N = 3$, let the dragons be called A, B, and C. After your announcement, C will think to himself, "If I do *not* have green eyes, then A and B can see that I don't, so as far as they are concerned, they can use the reasoning for the $N = 2$ situation, in which case they will both turn into sparrows on the second midnight." Therefore, if A and B do not turn into sparrows on the second midnight, then on the third day C will conclude that he himself must have green eyes, and so he will turn into a sparrow on the third midnight. The same thought process will occur for A and B, so they will all turn into sparrows on the third midnight. The pattern now seems clear.

Claim: *Consider N dragons, all of whom have green eyes. If you announce to all of them that at least one of them has green eyes, they will all turn into sparrows on the Nth midnight.*

Proof: We will prove this by induction. We will assume the result is true for N dragons, and then we will show that it is true for $N + 1$ dragons. We saw above that it holds for $N = 1, 2, 3$.

Consider $N + 1$ dragons, and pick one of them, called A. After your announcement, she will think to herself, "If I do *not* have green eyes, then the other N dragons can see that I don't, so as far as they are concerned, they can use the reasoning for the situation with N dragons, in which case they will all turn into sparrows on the Nth midnight." Therefore, if they do not all turn into sparrows on the Nth midnight, then on the $(N + 1)$st day A will conclude that she herself must have green eyes, and so she will turn into a sparrow on the $(N + 1)$st midnight. The same thought process will occur for the other N dragons, so they will all turn into sparrows on the $(N + 1)$st midnight. ■

Hence, in our problem all one hundred dragons will turn into sparrows on the 100th midnight.

Although we've solved the problem, you may be troubled by the fact that your seemingly useless information did indeed have major consequences. How could this be, when surely all the dragons already knew what you told them? Did you really give them new information? The answer is "yes". Let's see what this new information is.

Consider the case $N = 1$. Here it is clear that you provided new information, since you essentially told the one dragon that he has green eyes. But for the cases $N \geq 2$, the new information is slightly more subtle.

Consider the case $N = 2$. Prior to your announcement, A knows that B has green eyes, and B knows that A has green eyes. That is the extent of the knowledge, and they can't conclude anything else from it. But after you tell them that at least one of them has green eyes, then A knows *two* things: He knows that B has green eyes, *and* he knows that B knows that there is at least one dragon with green eyes (because A knows that B heard your information). B gains a similar second piece of information. This second piece of information is critical, as we saw above in the reasoning for the $N = 2$ case.

Consider the case $N = 3$. A knows that B green eyes, and he also knows that B knows that there is at least one dragon with greens eyes (because A can see that B can see C). So the two bits of information in the $N = 2$ case above are already known before you speak. What new information is gained after you speak? Only after you speak is it true that A knows that B knows that C knows that there is at least one dragon with green eyes.

The analogous result holds for a general number N. There is no paradox here. Information *is* gained by your speaking. More information is added to the world than the information you gave.[1] And it turns out, as seen in the proof of Claim 1, that the new information is indeed enough to allow all the dragons to eventually figure out their eye color.

To sum up: Before you make your announcement, the following statement is true for N dragons: A_1 knows that A_2 knows that A_3 knows that $\ldots$ that A_{N-2} knows that A_{N-1} knows that there is at least one dragon with green eyes. This is true because A_{N-1} can see A_N; and A_{N-2} can see that A_{N-1} can see A_N; and so on, until lastly A_1 can see that A_2 can see that $\ldots$ that A_{N-1} can see A_N. The same result holds, of course, for any group of $N - 1$ dragons. The point is that it is only after you make your announcement that the chain is extended the final step to the Nth dragon. The fact that the Nth dragon heard your statement is critical to the truth of this complete chain.

So, in the end, it turns out to be of great importance how far the chain, "A knows that B knows that C knows that $\ldots$" goes.

Note that if one of the dragons missed your farewell announcement (which was "At least one the 100 dragons on this island has green eyes"), then they will all happily remain dragons throughout the ages.

[1]For example, A knows that you made your statement *while* stepping onto your boat *and* wearing a blue shirt. Or, more relevantly, A knows that you made your statement in front of all the other dragons. In short, it's not just what you say; it's how you say it.