Solution to problem 7.55 in Ross's Probability.

A flock of N ducks goes by; the value of N is a random variable distributed following a Poisson distribution with expected value of 6. Ten hunters each randomly pick a duck in the flock to shoot at (independently and with uniform probability). Each hunter has a 60% probability of hitting his or her duck (independent of any other hunter or any previous flock). Each hit duck is killed and falls. What is the expected number of fallen ducks?

Before we begin, let's note a few "reasonability" tests we can place on our results later. The bigger the flock, the greater the number of expected fallen ducks since the probability of a hunter aiming at an already-dead duck decreases. Indeed, for a very large flock we would be very unlikely to have multiple hunters firing at any one duck; each hunter separately can expect to kill 0.6 ducks, so since expectations are additive we might expect about 6 ducks killed by the hunters altogether. At the other extreme, if $N = 0$ (and yes, that is allowed by the statement of the problem), there are no dead ducks, while if $N = 1$ every hunter aims for the one duck, which is then dead with probability $1 - (0.4)^{10}$, that is, we expect 0.999895 dead ducks in that case. With flocks of the "expected" size of 6, we might observe that there are many ways of matching hunters and ducks, but we might imagine that something like this might be typical: one duck not targeted, one duck targeted by 1 hunter, three ducks targeted by 2 hunters each, and one duck targeted by 3 hunters. Then the ducks' chances of survival are, respectively, 100%, 40%, 16%, and 6.4%, thus contributing $0 + .6 + 3 \times .84 + .936 = 4.056$ fallen ducks. In short, we learn (1) the expected number of fallen ducks depends on the size of the flock, and (2) a reasonable range for the answer to the question is perhaps around 4.

So let's set this up. For a sample space S, I would imagine the outcomes would each be the result of one flock passing before the hunters and getting fired upon. What I would record of each such outcome would be a picture which pairs 10 hunters with N birds, some of the pairings red (a hit) and the rest blue (a miss). So it's actually an infinite set of outcomes, not all of them equally likely. There are some random variables of interest: $N(s)$ is the number of ducks in a flock s; $X(s)$ is the number of ducks killed in outcome s; $H_j(s)$ is the number (0 or 1) of ducks killed by hunter j (I used this RV implicitly in my large-N analysis above); etc.

If I take the trouble to number the individual ducks within each flock, then also relevant for me would be $D_i(s)$, the number of times duck i dies in outcome s; this is obviously either 1 or 0, and the key variable is then $X = D_1 + D_2 + D_3 + \ldots$. You might object that, say, $D_3(s)$ is not well defined when s is a flock with only two birds, but that's OK: the whole point here is that we're going to compute our expectation by conditioning on N, the size of the flock. That is,

$$E[X] = \sum_{n=0}^{\infty} E[X|N = n] \cdot Pr(N = n)$$

where we have already calculated a few of those conditional expectations: 0, when $n = 0$; 0.9999 when $n = 1$; around 4 when $n = 6$; and increasing to 6 as $n \to \infty$.

Now, if we know $N = n$, then the random variables $D_1, D_2, \ldots, D_n$ all make sense and their sum is X, so that $E[X|N = n] = \sum_{i=1}^{n} E[D_i|N = n]$ by linearity. Of course by symmetry each of these summands is the same (there's no reason to assume any duck is more or less likely to be targeted, or to be hit when targeted, than say duck number 3) so that $E[X|N = n] = n\,E[D_3|N = n] = n\,Pr(D_3 = 1|N = n)$.

So just what is the probability that any particular duck i dies, knowing that there are n ducks in the flock? Well, of all the outcomes in the sample space that do have n ducks in them, hunter j aims at duck i in $1/n$ of the cases (since s/he will fire at each duck with equal probability) and kills this duck $0.6/n$ of the time, i.e. hunter j spares duck i with probability $1 - 0.6/n$. Thus ALL the hunters will spare duck i with probability $(1 - 0.6/n)^{10}$, and hence duck i dies with probability $1 - (1 - 0.6/n)^{10}$. Accordingly, I find our conditional expectation to be

$$E[X|N = n] = n\left(1 - \left(1 - \frac{0.6}{n}\right)^{10}\right)$$

(This is exactly what I computed earlier when $n = 1$, and it is easily seen to be an increasing function of n, as I expected. For large n, $0.6/n$ is small, and so we can expand the tenth power using just the first terms of the binomial expansion: $(1 - 0.6/n)^{10} = 1 - 6/n + 16.2/n^2 - \ldots$ making $E[X|N = n] = 6 - 16.2/n + 25.92/n^2 - \ldots$, again matching our reasonability checks. And finally, although we considered only one configuration of an outcome when $n = 6$, we got a value quite close to the true expectation of $E[X|N = 6] = 6\,(1 - (1 - 0.6/6)^{10}) = 3.908$)

Finally, then, we can combine this conditional expectation with the given distribution of flock size: we know

$$E[X] = \sum_{n=0}^{\infty} E[X|N = n] \cdot Pr(N = n)$$

but now since $E[X|N = n] = n\,(1 - (1 - 0.6/n)^{10})$ and $Pr(N = n) = e^{-6}6^n/n!$, we have

$$E[X] = \sum_{n=0}^{\infty} n\,(1 - (1 - 0.6/n)^{10}) \cdot e^{-6}6^n/n!$$

I don't know a way to simplify this sum but only a few terms add a non-negligible contribution to the sum (because the conditional expectations are all less than 6 and the tail probabilities are so small). I work out the sum to be approximately 3.69889875.