What is the distribution of unpecked chickens?

Jimmy Jin

July 15, 2017

I came across the following problem in a *New York Times* article about the 2017 Mathcounts National Competition.

> In a barn, 100 chicks sit peacefully in a circle. Suddenly, each chick randomly pecks the chick immediately to its left or right. What is the expected number of unpecked chicks?

This was the the final question in the final round of the competition and the winner, 13-year old Luke Robitaille, got it within a second.

Later, professor Steven Strogatz of Cornell tweeted a follow-up question by a fellow math enthusiast:

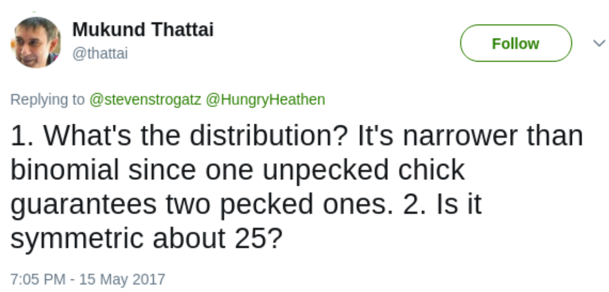
Mukund Thattai
@thattai
Follow

Replying to @stevenstrogatz @HungryHeathen

1. What's the distribution? It's narrower than binomial since one unpecked chick guarantees two pecked ones. 2. Is it symmetric about 25?

7:05 PM - 15 May 2017

And this problem has tormented me for hours. Finally, through the help of Ed Dewey and Zachary Abel though, we have an answer.

The expected number of unpecked chicks

Here's the easier problem. The answer is 25. The reasoning goes like this.

(1) Any fixed chicken goes unpecked if and only if both of its chicken neighbors peck **away** from it.

(2) This happens with probability 1/4.

(3) Therefore since each chicken has 1/4 chance of being unpecked and there are 100 of them, we can expect $25 = 100 \times 1/4$ to be unpecked on average.

Although this problem is easy, it's valuable because it highlights the power of one of the most taken-for-granted properties in probability: linearity of expectation.

Fact. For any random variables $X_1, X_2, \ldots X_n$,

$$\mathbb{E}(X_1 + \ldots + X_n) = \mathbb{E}X_1 + \ldots + \mathbb{E}X_n$$

Crucially, this theorem says nothing in it about the **dependence** between the random variables $X_1, \ldots, X_n$.

In step (3) above, we applied the "1/4-chance-of-remaining-unpecked" analysis to each of the 100 chicks in the circle. But there's a problem. For chick #1 to remain unpecked, its neighbors (chicks #100 and #2) must peck away from chick #1. But chicks #100 and #2 are also neighbors of other chicks as well. Going all the way around the circle, we see that there's a ton of dependence. It's a pain in the ass.

Linearity of expectation lets us ignore that when we're looking at averages. The reason why this property is so taken for granted is that it seems ridiculously obvious when the situation involves independence. Nobody ever bats an eyelid when we use it to show that the expected number of heads is 25 when flipping a coin 100 times with $p = 1/4$ of coming up heads.

Indeed, when dependence is involved, that's when linearity of expectation is at its most powerful. Anyway, moving on...

The distribution of unpecked chicks

Now here's the hard one. The answer depends on whether there are an odd or even number of chicks.

$$\mathbb{P}(\text{exactly } k \text{ unpecked chicks}) = \begin{cases} 2\binom{n}{2k} \Big/ 2^n & \text{if } n \text{ odd} \\ \left[2\binom{n}{2k} + 2(-1)^k\binom{n/2}{k}\right] \Big/ 2^n & \text{if } n \text{ even} \end{cases} \quad (1)$$

I first started thinking about this problem because I saw Aaron Clauset propose a Binomial distribution as the answer, maybe Binomial(100, 1/4). But there are a few issues.

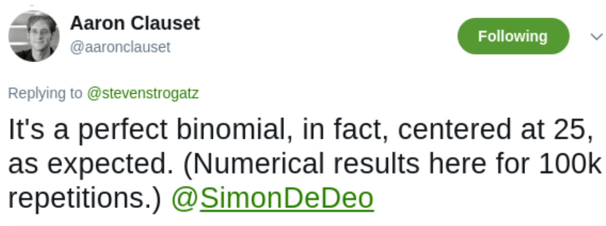

1. First of all, $\mathbb{P}(\geq 50 \text{ unpecked chicks}) = 0$, so it can't be Binomial(100, 1/4).
2. Secondly, even if you switch to Binomial(50, 1/2), then you would expect to be able to cast it as a sum of 50 independent trials with probability 1/2 of success. And as we saw, there's no obvious way to cut up the circle of 100 chickens into pairs of independent chunks. The dependency in this problem suggests something other than Binomial.

It turns out that Prof. Clauset just had a bug in his code. Once he fixed it, it found that the distribution was actually a lot narrower than Binomial:

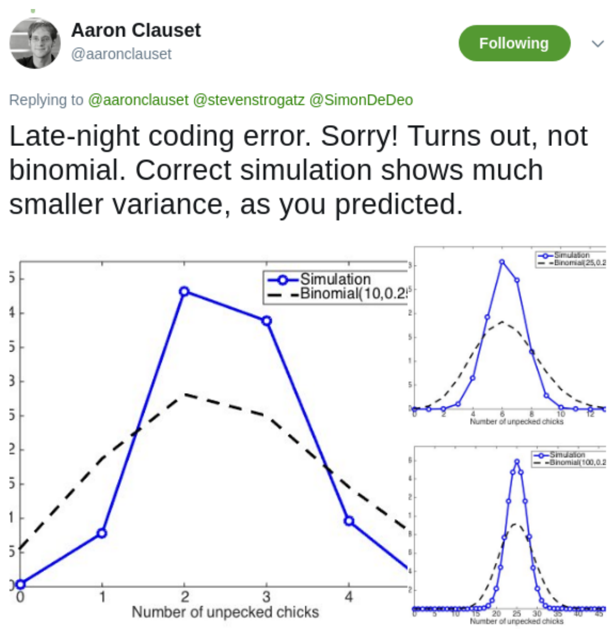

So. . . what is it?

Since the chicks sit in a circle and each peck only to one side or the other, you can view the setup as a directed ring graph where the chicks are directed *edges*, not nodes.

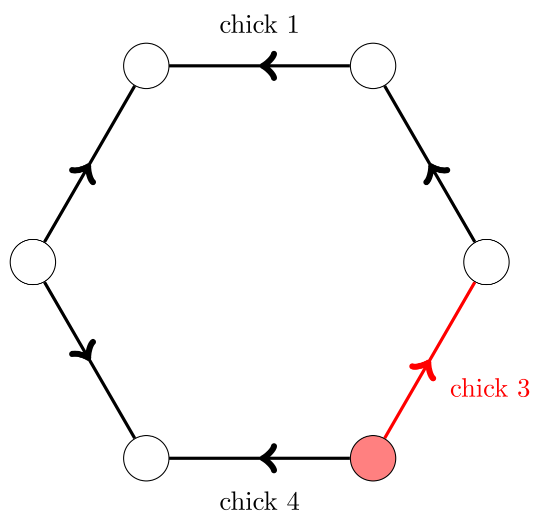

Figure 1: A realization of pecks in a circle of $n = 6$ chicks. Chick 1 pecks left towards its neighbor Chick 6. Chick 6 pecks right back at Chick 1. The lone red circle and edge are associated with the single unpecked chick, Chick 3.

If you look at it this way, it's easy to see that a unpecked chick occurs if and only if you have a triple of adjacent edges which have a very certain configuration. This is the configuration which represents the situation of *a chicken's two neighbors pecking outwards while the chicken itself pecks in either direction.*

Call this configuration an **out-triple**. The single out-triple corresponding to the unpecked chicken in Figure 1 is marked below.

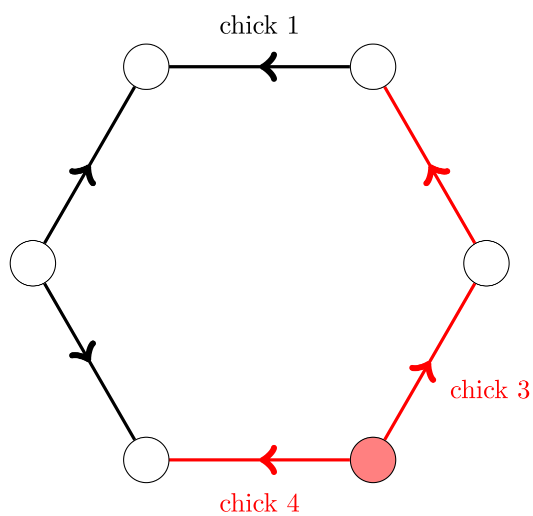

Figure 2: The realization from Figure 1 with the out-triple highlighted in red.

The easy part is noting that:

> **Proposition.** For a fixed number of chicks n, the number of unpecked chicks has the same distribution as the number of out-triples in a directed ring[1] on n vertices, where each edge is directed left or right with equal probability and independently of all the other edges.

Therefore we will be done if we can just find the distribution of out-triples in these random directed rings. Since every configuration of edges in these random directed rings is equally likely and there are 2^n of them total, we just need to count the number of configurations that have exactly k out-triples and divide by 2^n to get the final answer.

Unfortunately, counting these triple configurations is annoying. This is because two out-triples can "nest" within each other, so it looks like we might need to keep track of not only out-triples but also **left** and **right** out-triples, corresponding to the pecking direction of the central (unpecked) chick.

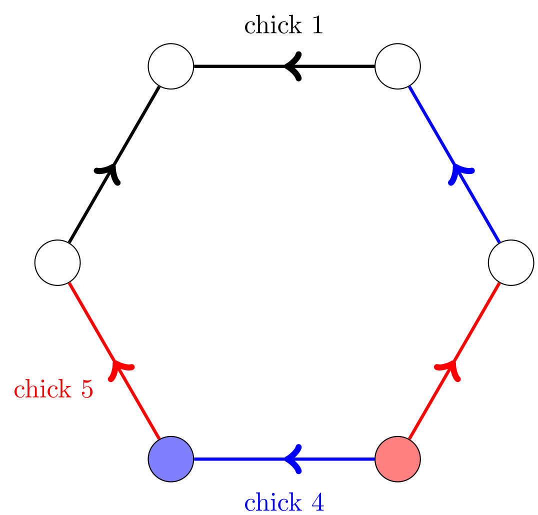

Figure 3: The realization from Figure 1, except now Chick 5 pecks to the left. This creates two nested triples and therefore two unpecked chicks right next to each other.

Fortunately, there's a slightly easier way to deal with this.

First, count the out-*pairs*

I provided an earlier solution to this problem that was wrong, because I thought unpecked chicks corresponded to adjacent *pairs* of edges which pointed away from each other. I called this an **out-pair**.

[1] I call it a ring because, strictly speaking, a directed cycle has all the directions pointing the same way.

Curiously, the final solution to this turned out to be incorrect for even n, but correct for odd n. What gives? As it turns out, the entire problem about out-triples can be reduced to one about out-pairs. So first, let's start by counting configurations with out-pairs:

> **Lemma.** The number of configurations on a directed ring with n vertices that have exactly k out pairs is $2\binom{n}{2k}$.

I learned the following elegant argument for this fact from my friend Ed Dewey:

1. The entire edge configuration is determined by the in-pairs and out-pairs.
2. Also, in-pairs must alternate with out-pairs, i.e. between any two in-pairs there is an out-pair and vice versa. Additionally this means that there is one in-pair for each out-pair.
3. Let S be the union of all the vertices at the center of an in-pair or an out-pair. Then the whole directed ring is determined once we know S and whether a single element of S belongs to an in-pair or an out-pair.
4. S has size $2k$ since we require k out-pairs, and every out-pair has a matching in-pair.

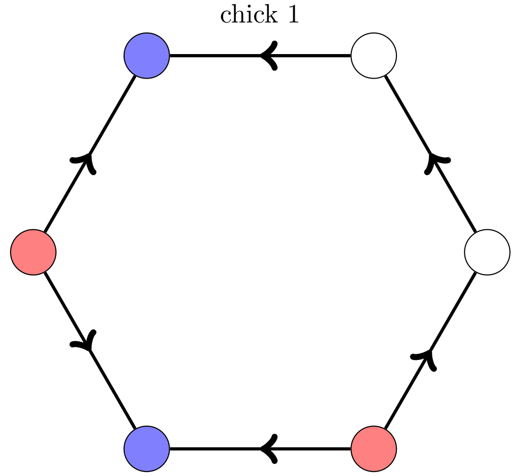

Figure 4: The realization from Figure 1. Red vertices belong to out-pairs. Blue vertices belong to in-pairs. Red and blue must alternate, and S is the union of the red and blue vertices. The paths adjacent to the white vertices (not belonging to either an out-pair nor an in-pair) are completely determined by the configuration of the red and blue vertices.

From out-pairs to out-triples

Here's the key observation which I learned from Zachary Abel. When you back up and look at the ring of chicks, *you can break it up by looking at every other chick.*

The reason is simple. Whether or not a chick is pecked depends solely on its two neighbors. And therefore, the event of whether or not a chick is pecked can only influence *the event of whether another chick is pecked* if and only if that other chick is within a multiple of 2 away from it.

Be careful here. Obviously, if we know that a chick is pecked, we know something about the orientation of the pecking direction of its neighbors. However, we do not know anything about whether or not its neighbor is *pecked*. That is the crux of the matter.

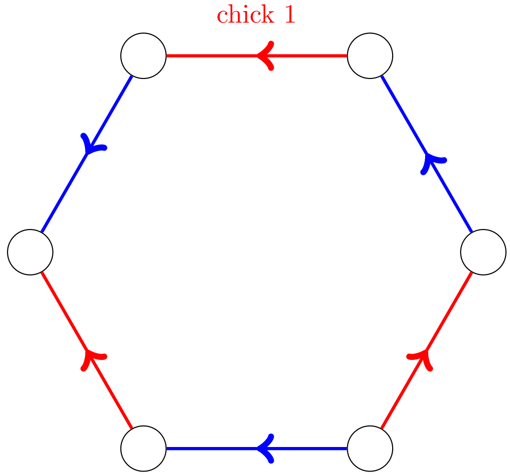

Figure 5: The realization from Figure 3 with edges colored according to whether the chick is even or odd. Knowing whether an even chick is pecked does not tell us whether an odd chick is pecked, at least when n is even.

To put it another way, even though we've observed that there is a lot of dependence in this ring, Zachary noticed (loosely speaking) that the dependency between **the events of whether or not the chicks are pecked** is contrained to influencing only between every other chick.

Our strategy will therefore be to split the ring of chicks into two separate rings: one for odd chicks and one for even chicks. And because there is independence between the two, we can analyze them separately.

But there's a wrinkle though. When there's an odd number of chicks in total, going around to every other chick eventually covers all the chicks. Therefore we'll need to split into two cases:

1. When n is even, we split the ring into two separate rings: one for the even-numbered chicks and one for the odd-numbered chicks. Call these the **odd and even rings**.
2. When n is odd, the even and odd rings coincide and we'll just deal with the original ring.

Now looking at the odd and even rings, we can clearly see that counting out-triples in the original ring is equivalent to counting out-pairs in the odd and even rings.

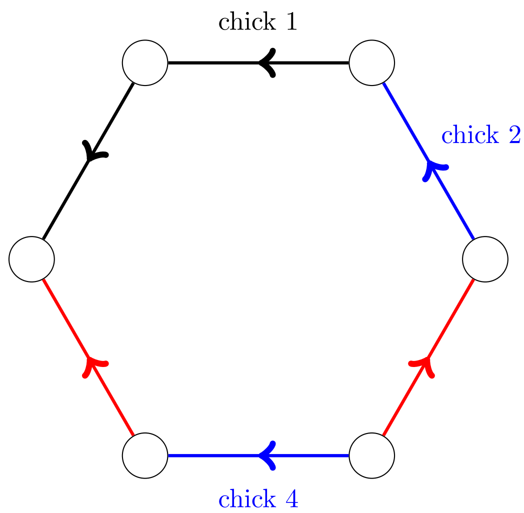

Figure 6: The realization from Figure 3 with the two nested out-triples.

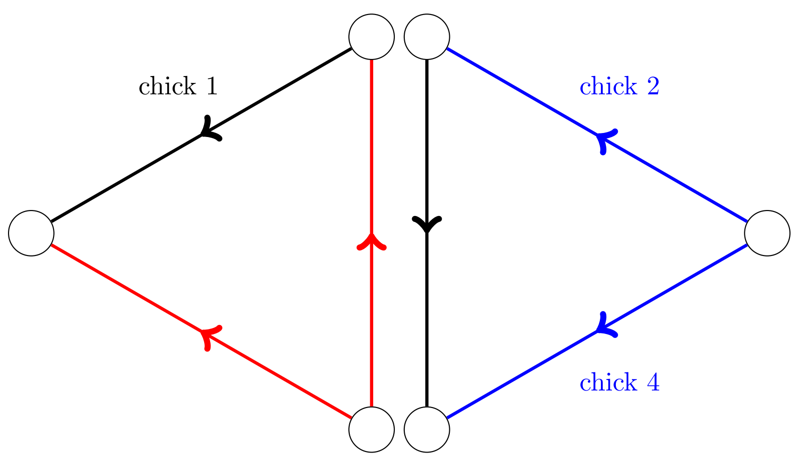

Figure 7: The realization from Figure 3 with the two nested out-triples, separated into odd and even rings. The two nested out-triples now manifest as two seprated out-pairs which appear independently of each other.

It therefore turns out that the total number of configurations with k out-triples is the number of configurations with k out-pairs in the odd and even rings. What is this number?

1. **For odd** n: this is just the number of out-pairs in the entire ring, since the odd and even rings are just the entire ring. That's

$$2\binom{n}{2k}$$

2. **For even** n: we have to do a bit more. But using this decomposition, we can count it like this.

   Let $A_n(k)$ be the number of configurations of n chicks that have exactly k out-**pairs**. Then the number of configurations of n chicks with exactly k out-triples denoted $B_n(k)$ is:

$$B_n(k) = A_{n/2}(0) \cdot A_{n/2}(k) + A_{n/2}(1) \cdot A_{n/2}(k-1) + \ldots + A_{n/2}(k) \cdot A_{n/2}(0)$$

   In English, the number of configurations with k out-triples is equal to the number of configurations with 0 out-pairs in the odd ring times the number of configurations with k out-pairs in the even ring, plus the number of configurations with 1 out-pairs in the odd ring times the number of configurations with $k-1$ out-pairs in the even ring, etc...

   When you add this up, you get our answer for even n:

$$2\binom{n}{2k} + 2(-1)^k\binom{n/2}{k}$$

   The fancier way to look at it is that we are taking the out-pair count for a single ring of size $n/2$ and convolving it with itself. Doing this using generating functions as Zachary suggested, the generating function for a single $n/2$ ring is

$$R(x) = (1+\sqrt{x})^{n/2} + (1-\sqrt{x})^{n/2}$$

   Therefore the generating function for what we want is $R^2(x)$, and extracting the coefficient on x^k in the series expansion of $R^2(x)$ is exactly the answer given in the display above.

Taking these quantities (for odd n and even n) over 2^n, the total number of configurations of the directed ring on n vertices, gives the distribution presented at the top.

Phew. Glad we put that to rest. Thanks for your help, Ed Dewey and Zachary Abel!