Nandor's solution to the Bone Man Problem

Introduction

First, you'll have to exclude the long-windedness. I'm a physicist by training and I teach High School Math. I therefore have a penchant for lots of data and length explanations.

I've worked on this problem a number of different ways, but each of them makes sense, so here goes (I unfortunately have no idea what the Oz book
(http://sprott.physics.wisc.edu/pickover/mathozad.html)
says on the matter, since I haven't read it).

If you want to skip the boring stuff and skip to the punch line, scroll down to the bottom. You'll see some pretty Monte Carlo simulation results.

Here is the original problem statement:
Dorothy and Dr. Oz peer into a deep hole in the ground. The bone man comes closer and opens and closes his mouth spasmodically. "In the pit," he says, "are 10,000 leg bones. I have cracked each bone **at random** into two pieces by throwing them against a rock. What do you think is the average ratio of the length of the long piece to the length of the short piece for each time I crack a bone? You can reason from a purely theoretical standpoint. If you cannot find the solution within two days, I will add your leg bone to the pit."

The key statement, as we all know, is "What do you think is the average ratio of the length of the long piece to the length of the short piece for each time I crack a bone." Note that it would be impossible to give an answer that would be correct all the time, since all of the bones, if we were unlucky, could break 99% of the way to one end and the average (whether you're talking mean, median, or mode) would be 99 (a 99:1 ratio). What we are trying to do, then, is salvage whatever we can out of the situation by picking a more likely average ratio than 99 (if a more likely average ratio even exists).

Here are my answers for the Bone Man problem:

1) If by "average" the Bone Man means "mode," there is no good answer since there will be an essentially flat distribution in the pile - every ratio of the 10,000 bones could well be different, so there is no mode for his 10,000 bones.

2) If by "average" the Bone Man means "median," then the answer is 3, since half the time the break will occur between 75% and 100% of the length of the bone and the other half of the time the break occurs between 50% and 75% of the length of the bone. The median value is thus 75% of the length; the short piece will then be 25% of the length and the ratio is 3:1.

3) If by "average" the Bone Man means "mean," which is the mostly likely interpretation, then there is a great solution, but it COMPLETELY depends on the increment of measurement. Once we nail down an increment of measurement, we can get a handle on the surprising answer. FROM HERE ON OUT, IF I SAY "AVERAGE" IT INDICATES THE MEAN.

The 3 Keys

The main key to this problem is to ask the Bone Man the proper question before answering. You MUST know how good his ruler is. Assuming that all bones start off approximately the same length, let's call it 1 unit long. We MUST know what the accuracy of the ruler is. Can the Bone Man measure to 1/100th? 1/10th? 1/100,000,000th? We'll see later why this matters.

Another key is to realize that the question REALLY is "what is the most likely average ratio of the long piece to the short piece?" We can see this by understanding that his bone pile, if he can measure to 1/100th, could easily be 99, if every bone were unlucky enough to break 0.99 of the way toward one end. Let's think about this for a moment. 99 is a VALID average ratio, but it certainly is not a likely average. In fact, the probability of this being the average (when we measure to the nearest 1/100th) is

$$\left(\frac{1}{49}\right)^{10000} = 1.09 \cdot 10^{-16902}.$$

I get this since there are 49 different ratios, from .51/.49 through .99/.01 (there can't be 50/50 since we are told that there is always a long piece and a short piece, and there can't be 100/0 since such a bone would not have broken), and there are 10,000 bones in the yard. We then have 10,000 places to fill with 49 different choices, and repeats are allowed.

So the total number of different ratios is $49^{10000} = 9.137 \cdot 10^{16901}$, and thus the probability of getting all 99s is only $1.09 \cdot 10^{-16902}$, which is, incidentally mathematically zero since we tend to call something zero if it is less than 10^{-100}.

So 99 as an AVERAGE RATIO is unlikely. The question, then, is what AVERAGE RATIO is the most likely?

The last key before we start actually working on the problem is to ask the Bone Man how close you have to be. If the actual average ratio in his pile is 7.24534 and you say 7.245, are you correct? What if you say 7.2? Or 7? Or 6.8? It's best to know what you're getting into!

Starting The Problem

So, armed with these three keys, let's start with a simple problem that is similar. Let's say that the Bone Man's ruler can only measure to 1/10th. Thus there are only four breaks that he can measure: 0.6, 0.7, 0.8, and 0.9. The ratios associated with these four breaks are 1.5, 2.333333..., 4, and 9.

1 Bone
Let's start with a 1-bone boneyard. The average will thus simply be the ratio of that single bone in the yard. Since each ratio has an equal probability of occurring, you have a 1/4 probability of being correct no matter which average you choose. Your best bet is to choose one of the four possible average ratios and hope for the best.

2 Bones

Let's now look at a 2-bone boneyard. Now there are more averages possible: $\left(\!\!\binom{4}{2}\!\!\right) = \binom{5}{2} = 10$.

The (( )) is called the multi-choose function, and it is similar to the ( ) choose function, except we must use multi-choose when repeats are allowed. The reason that

$$\left(\!\!\binom{4}{2}\!\!\right) = \binom{5}{2}$$

is that a simple formula relates multi-choose and choose:

$$\left(\!\!\binom{N}{k}\!\!\right) = \binom{N+k-1}{k}.$$

The total number of arrangements are 16, so all of the probabilities of those averages occurring will be multiples of 1/16. Note that there must be 10 different columns.

Two ratios in the yard	Average Ratio	Probability of Occurring
1.5 and 1.5	1.5	1/16
1.5 and 2.3333333	1.91666666	1/8
1.5 and 4	2.75	1/8
1.5 and 9	5.25	1/8
2.3333333 and 2.3333333	2.33333333	1/16
2.3333333 and 4	3.16666666	1/8
2.3333333 and 9	5.66666666	1/8
4 and 4	4	1/16
4 and 9	6.5	1/8
9 and 9	9	1/16

The reason that some of the ratios have 1/8 probability is that there are two ways for them to occur (for instance, the two ratios in the yard could be 1.5 and 4, or they could be 4 and 1.5). The other probabilities are only 1/16 since there is only one way to have both bones be the same. Note that another way to think about this is to pretend that there are 16 different bone yards, each with only two bones in them. Assuming everything is kosher, we can EXPECT (it isn't guaranteed, of course) that two of them will have an average ratio of 6.5 and two of them will have an average ratio 5.25 and one of them will have an average ratio of 1.5, and so on.

So in this case, what should we tell the Bone Man? Well, we shouldn't tell him "the average ratio will be 2.3333333," since there is only a 1/16 chance of that happening. Instead, assuming we must be exact, we should choose one of the six different averages that give us a 1/8 chance of being correct. Choosing another answer doesn't make sense.

Here's where it might matter, however, how close we have to be. I arranged the table above in an order that makes sense for ordering the possible outcomes. If I instead arrange it by increasing numerical order of Average Ratio, the table looks like this:

Two ratios found	Average Ratio	Probability of Occurring
1.5 and 1.5	1.5	1/16
1.5 and 2.3333333	1.91666666	1/8
2.3333333 and 2.3333333	2.33333333	1/16
1.5 and 4	2.75	1/8
2.3333333 and 4	3.16666666	1/8
4 and 4	4	1/16
1.5 and 9	5.25	1/8
2.3333333 and 9	5.66666666	1/8
4 and 9	6.5	1/8
9 and 9	9	1/16

If we only need to be with 0.5 of the answer, then if we answered "2.1," we have a 3/16 chance of being correct. If we answer "1.9," we have a 1/4 chance of being correct. If we answer "2.75," we have a 5/16 chance of being correct. If the Bone Man allows us to be within 0.5 of the answer, then, we should answer somewhere between "2.66666" and "2.8333333." Any of these answers would get us to within 0.5. So you can see that the answer changes, depending on how accurate YOU have to be.

3 Bones
Continuing on, then, we now look at what happens where there are three bones in the boneyard.

There are 64 (4^3) different arrangements and $\left(\!\!\binom{4}{3}\!\!\right) = \binom{6}{3} = 20$ of them are distinct.

Three ratios found	Average Ratio	Probability of Occurring
1.5, 1.5, 1,5	1.5	1/64
1.5, 1.5, 2.333333	1.777777	3/64
1.5, 1.5, 4	2.33333333	3/64
1.5, 1.5, 9	4	3/64
2.33, 2.33, 2.33	2.33333333	1/64
2.33, 2.33, 1.5	2.05555555	3/64
2.33, 2.33, 4	2.88888888	3/64
2.33, 2.33, 9	4.55555555	3/64
4, 4, 4	4	1/64
4, 4, 1.5	3.16666666	3/64
4, 4, 2.33333	3.44444444	3/64
4, 4, 9	5.66666666	3/64
9, 9, 9	9	1/64
9,9,1.5	6.5	3/64
9, 9, 2.33333	6.77777777	3/64
9, 9, 4	7.33333333	3/64
1.5, 2.333, 4	2.61111111	6/64
1.5, 2.333, 9	4.27777777	6/64
1.5, 4, 9	4.83333333	6/64
2.3333, 4, 9	5.11111111	6/64

From a cursory glance at the table above, it appears that we could choose any of 16 different average ratios and have an equally good chance of being correct. But look again and you will see that some of the ratios are repeated, as a quirk of some of the averages being the same, even though they originate from different compositions of bone breaks. To have the best chance at being correct (3/32), we should choose to answer any of "2.611111," "4.277777," "4.833333," "5.1111111." If you have 0.5 leeway, you should guess a number between 4.61111 and 4.77777, giving yourself a 32.8% chance of being correct.

THIS IS AN IMPORTANT POINT! Not only do we have to look at the probability of each combination of bones, but we have to see which combinations give us the same ratios! The fact that some of them yield the same average ratio is the key to the final solution.

10,000 bones – or not?

Okay, so it seems like we could solve this problem (still for only measuring to the nearest 1/10th) for 10,000 bones with a computer. All we must do is determine the multinomial coefficients (the probabilities shown in the third columns above) for all of the

$$\left(\!\!\binom{4}{10{,}000}\!\!\right) = \binom{10{,}003}{10{,}000} = 166{,}766{,}685{,}001$$

possible ratios. After we do that, we need to combine the average ratios that are identical, adding their probabilities together. We then look at which average ratios have the best probabilities of being correct, and choosing one of those. Piece of cake. Of course, since there are over 100 billion of them, it may take some time.

Here's the deal. Using only those 4 breaking points (assuming the an measurement accuracy of 1/10th) it takes my computer about 3 minutes to determine all of the probabilities for only 100 bones. I'm not about to run through the calculation for 10,000 bones since the number of different combinations jumps from 4598126 to 166766685001. It would take my computer at least 75 days to run the calculation, assuming nothing slows it down, like running out of memory.

So this method is impossible to use, even for a measurement as crude as 1/10th. You can tell that this method won't work for determining the most likely average ratios for more accurate measurements. Instead I'll use a Monte Carlo simulation to find the best average ratio. Before I go on, though, I'll at least show what I got for those 100 bones.

This first graph below shows the raw probabilities (y-axis) of each possible average ratio (x-axis). Note that the average ratios only run from 1.5 through 9, as we already knew they should for our measurement accuracy (1/10th). Also note that for many of the average ratios between 3.5 and 5 seem to occur multiple times, since many averages ratios result from different combinations of bone breaks. The next step, then, is to add those duplicate average ratios together.

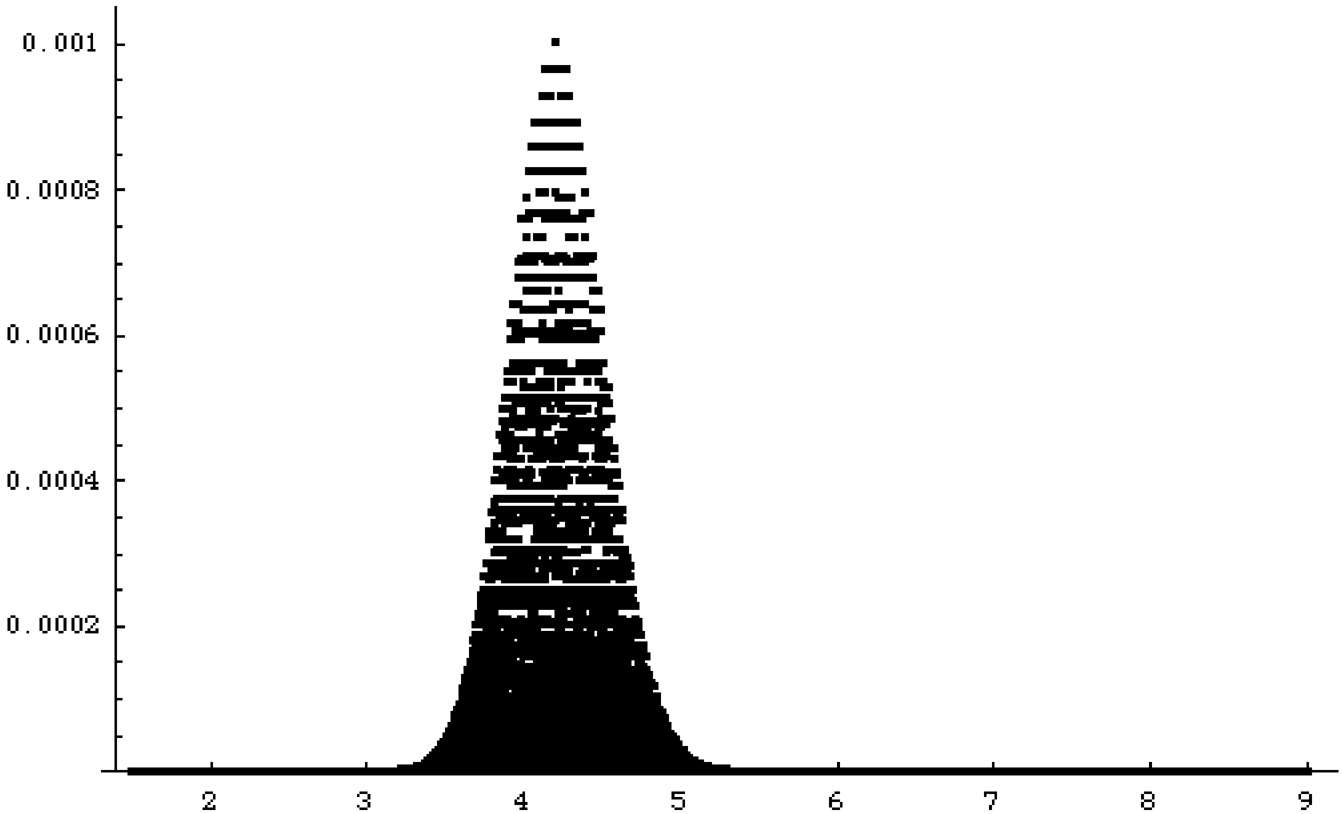

The next graph is the result you get when you add those all together.

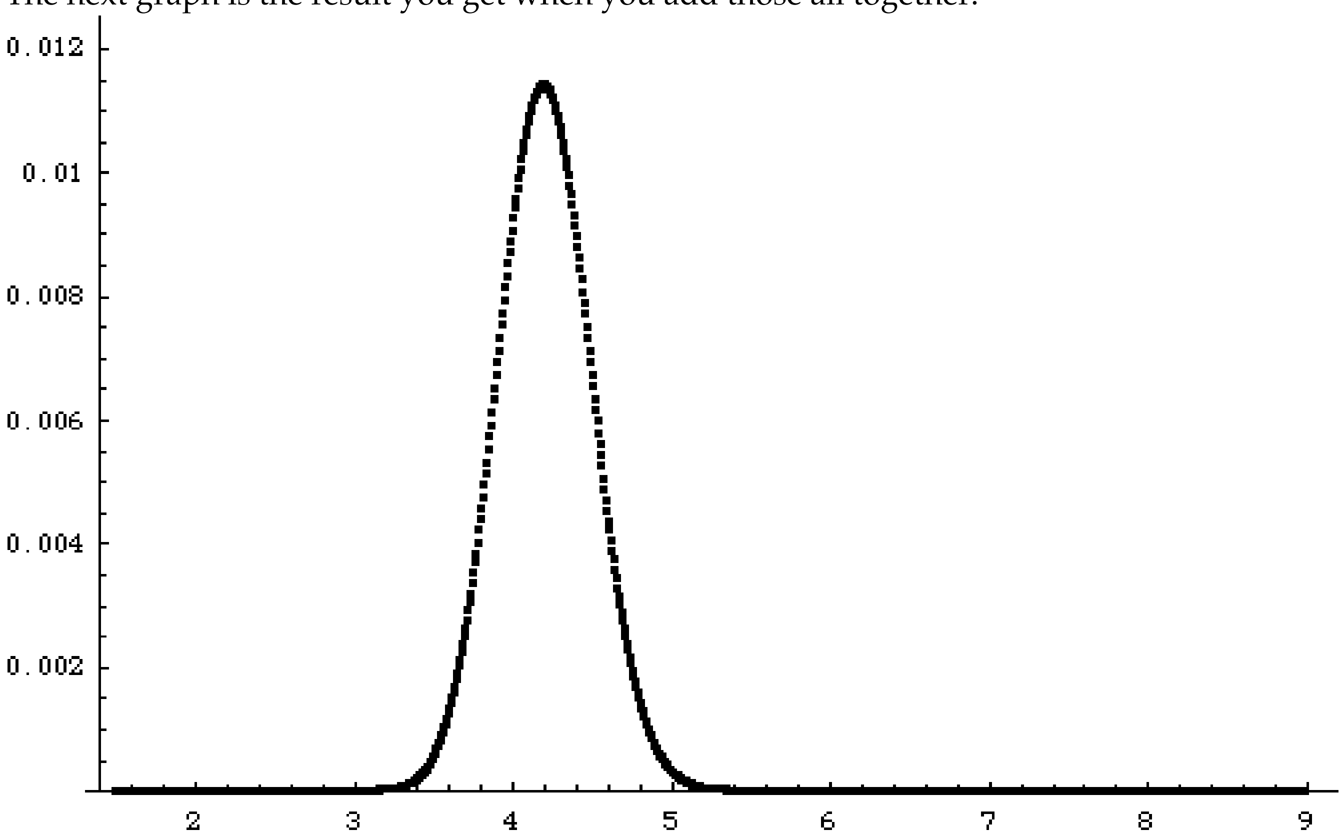

Remarkably, the result looks like your simple binomial distribution. As of yet, I don't have an explanation as to why it works out SO neatly. I wouldn't expect the average ratios to combine so nicely since such disparate numbers are being used to find those average ratios. However, the result is confirmed by my Monte Carlo simulations (see below), so there you have it.

Incidentally, since every bone ratio is equally likely, one might think that the average ratio with the highest probability would be THE one that is simply the average ratio of the four possible ratios (1.5 + 2.3333333 + 4 + 9)/4 = 4.20833333. In fact, this is CLOSE to the correct answer. There are actually two average ratios with slightly higher probabilities. The average ratio 4.2 has a probability of 0.01140627 of occurring. The average ratio of 4.19166666 has a probability of 0.01140617 of occurring. The average ratio of 4.20833333 "only" has a probability of 0.011397 of occurring. Yes, they're all about the same. But if my leg bone were at stake and I had to be exact, I'd have to say 4.2 and not 4.20833333. With any luck, the alien will be lenient and our answer will only have to be accurate to 0.5 or so. Adding up all of the probabilities from 3.7 to 4.7, we see that we would have a probability of matching HIS average ratio of 0.916977. Not bad.

But again, this is only for measurements of 1/10th. As stated above, then, we must find another way to solve the problem.

We could, as Mark as suggested on this discussion group, simply find the average of all of the possible ratios in the pile and use that. After all, it worked for measurements of 1/10th. This method DOES always work, as the simulations show, as long as the number of possible measurements is less than 10,000. So. If the Bone Man can only measure to the nearest 1/1000th, we're already done.

The formula to we can use, which is a "simplification" of our sum, is **equation 1:**

$$MostLikely = \frac{2 - M + 2 \cdot M \cdot \mathrm{EulerGamma} + 2 \cdot M \cdot \mathrm{PolyGamma}[0, \frac{M}{2}]}{M - 2},$$

where M is one over our precision (if we can measure to the 1/100th, M=100). Below is a plot of all of these values from M=10 to M=100,000 (actually, only the even ones – the odd ones give a slightly different-looking graph). EulerGamma, by the by, is equal to 0.5772156.... and is related to the remainder of an incomplete Harmonic Series.

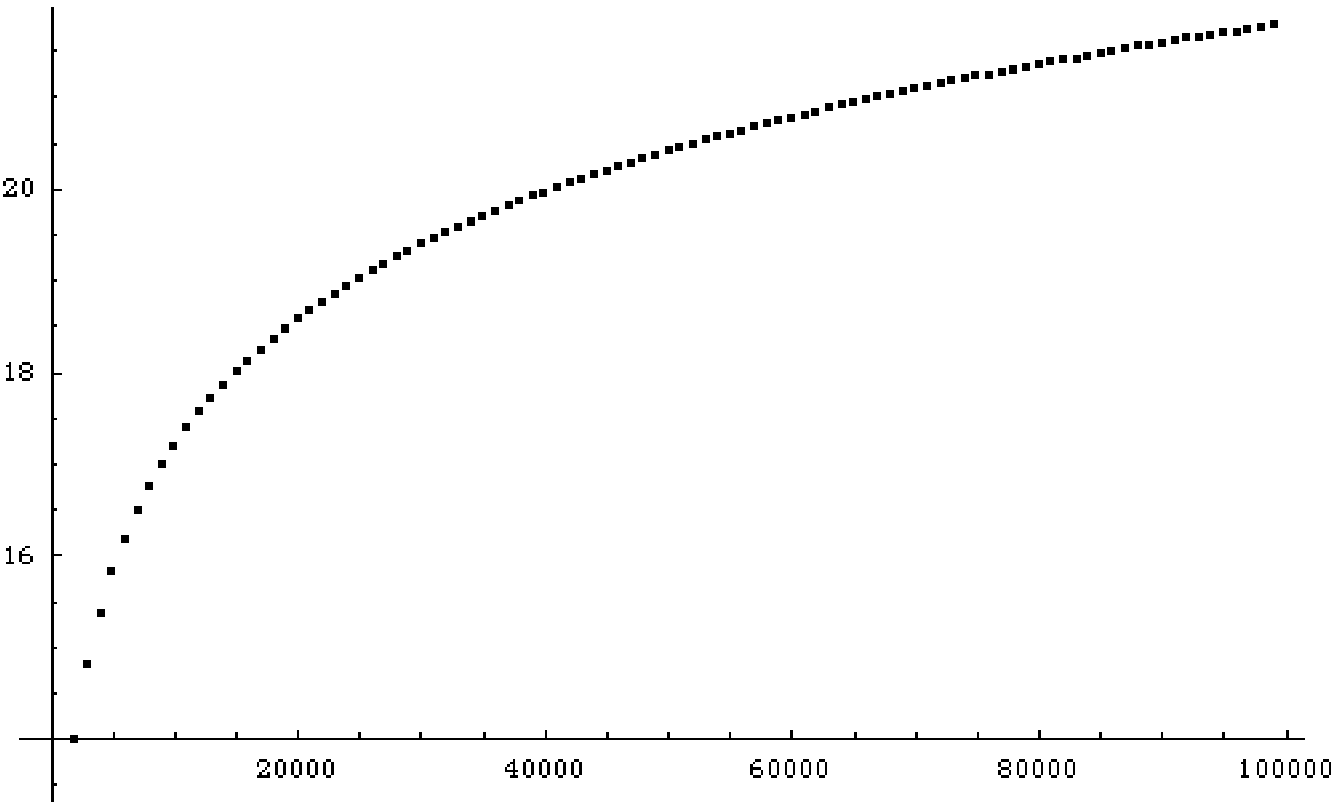

As you can see, the function is VERY slow. In fact, it looks remarkably like a Logarithm. In fact, if we plot it on a LogLinear plot (where the x-axis is logarithmic), the graph is almost a line (with a y-intercept of –EulerGamma!)

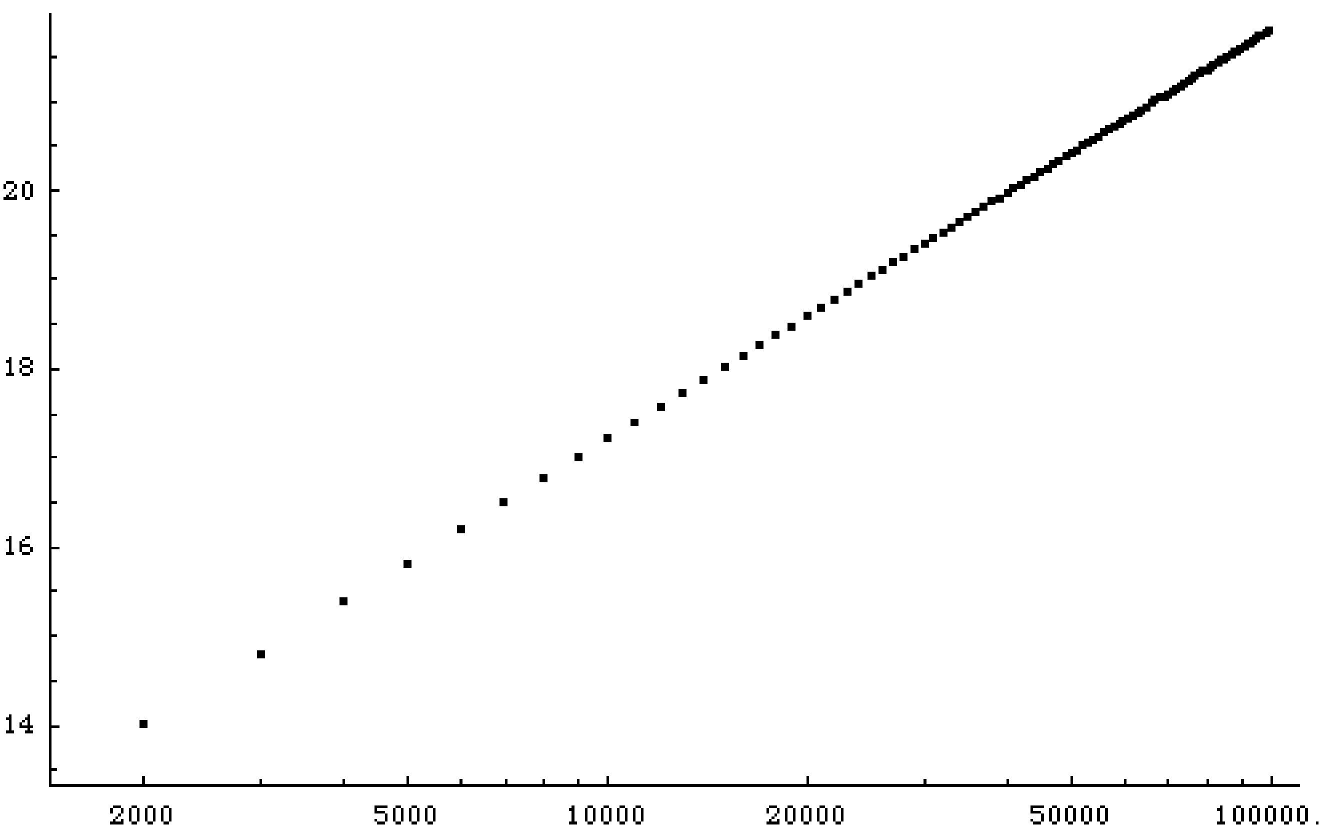

So, if the Bone Man says "I can measure to 1/458th," then you just have to plug 458 into the above formula and that WILL BE your best answer. It will give you the highest probability of success. No guarantees, but that's the most likely average ratio and your best chance at keeping your leg bone inside your leg. But that's not the end of the story, because although I have shown the graph above as going out to measurements of 1/100,000th, we can't trust the graph beyond measurements of 1/10,000th!

The Monte Carlo Simulations

Alright, enough stalling. On with the Monte Carlo simulations.

Here is the Mathematica script that I wrote for to simulate a pile of 10,000 bones if our measuring implement can measure to 1/10th of a unit (where all the bones are approximately 1 unit long):

```
<<Graphics`Graphics`
<<Statistics`DescriptiveStatistics`
breaks:=Table[N[Random[Integer,{6,9}]/10],{10000}]
ListOfMeanRatios={};
ListOfMedianRatios = {};
```

The first two lines load the necessary packages. "breaks" is a table of random breaks in the bone. Note that I only look at breaks that are 6/10 through 9/10. I don't care about 4/10 since I've already taken care of it with 6/100.

So now I have this list of 10,000 breaks. I've also initiated a couple of lists that I'll talk about next.

```
Do[{ratios = 1/(1/breaks - 1),
   ListOfMeanRatios = Append[ListOfMeanRatios, Mean[ratios]],
   ListOfMedianRatios = Append[ListOfMedianRatios, Median[ratios]],
   If[Mod[i, 100] == 0, Print[i]]}, {i, 50000}]
Histogram[ListOfMeanRatios, HistogramCategories -> 1000]
Histogram[ListOfMedianRatios, HistogramCategories -> 1000]
```

I'm running this simulation 50,000 times (think of it as 50,000 Bone Men). For each of those times, I have a boneyard of 10,000 bones. The first thing I do for each boneyard is to determine the ratio created by each break. The ratio is break/(1-break). The way Mathematica works, I can't mention "break" twice in one expression or else it generates two different 10,000-bone boneyards. So that's why it appears in the script as a simplified expression 1/((1/breaks) – 1).

All that happens, simply, is that I have calculated all of the ratios for those 10,000 bones. I then find the mean of those ratios and the median of those ratios. I then add that mean and that median to a list. Then I repeat. I repeat it 49,999 times!

The result is a 50,000-element list of mean ratios. I also have a 50,000-element list of median ratios. I then simply plot a histogram for each of them. The Print statement is so that I can keep track of the script's progress.

1/10th
Here is the histogram for a ruler that measures to 1/10th. The peak of the distribution, just like we saw above, is somewhere near 4.2. Just like we'd expect not only from what we saw above, but from the average of all of the possible ratios, (6/4 + 7/3 + 8/2 + 9/1)/4 = 4.20833333. Remember that what this histogram shows is a set of average ratios. Any of these average ratios below CAN happen, but obviously the average ratios near 4.2 happen more often than those average ratios near 9. Note the similar shape of this curve, to the bottom graph of the added probabilities on page 7 (which was only a calculation for 100 bones and not 10,000). It seems that we're on the right track both theoretically and simulation-wise based on all of our assumptions (evenly-distributed bone breaks, &c.)

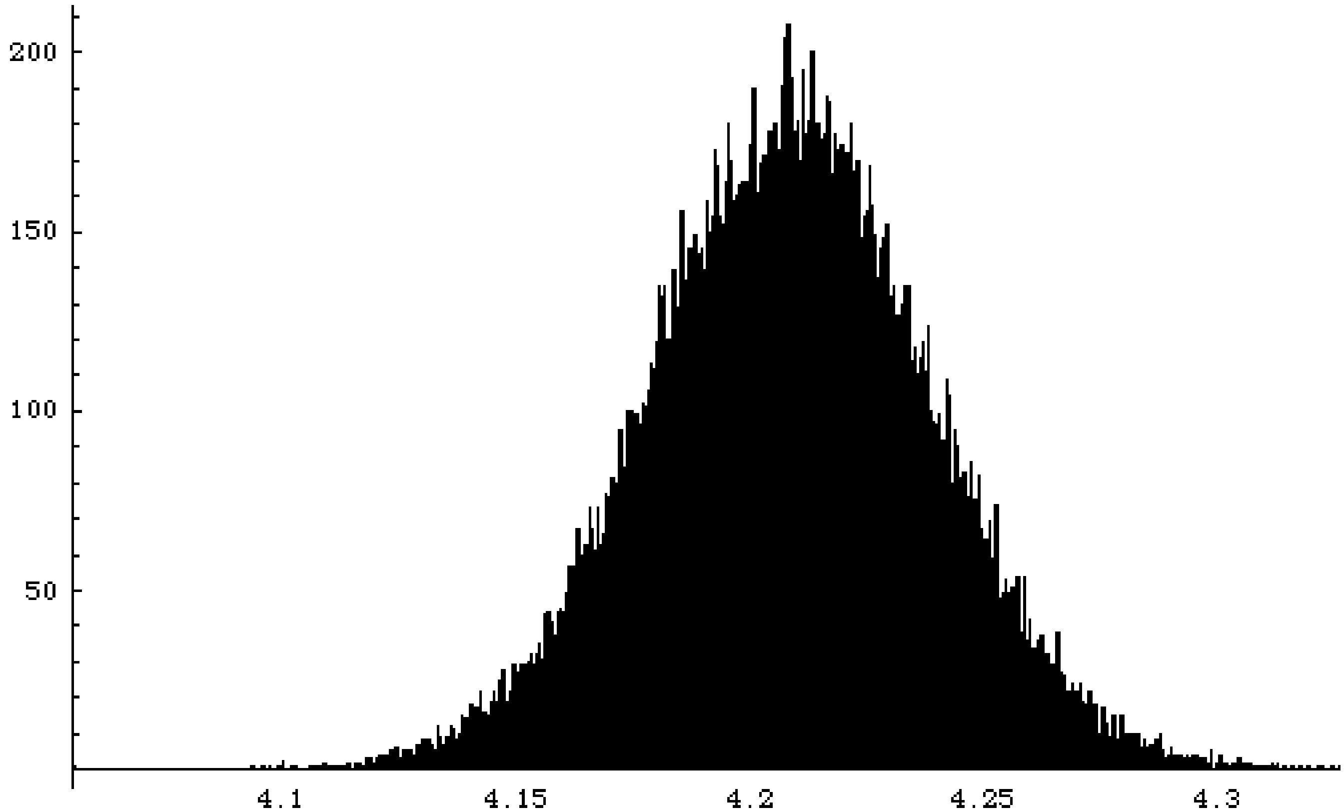

Distribution of Mean Ratios for measurements of 1/10th – 50,000 trials, bin size ~0.0003

1/100th
Next I ran the simulation as if our ruler could measure 1/100th of a unit.

```
<<Graphics`Graphics`
<<Statistics`DescriptiveStatistics`
breaks:=Table[N[Random[Integer,{51,99}]/100],{10000}]
ListOfMeanRatios={};
ListOfMedianRatios = {};
```

Not much of a difference here. To make a guess before we actually see the histogram, we would think that the most probable average ratio should occur at the same average of all of the ratios:

(51/49 + 52/48 + 53/47 + 54/46 + ... + 98/2 + 99/1)/49 = 8.141.

Below is the histogram of average ratios for the 50,000 different boneyards. Notice that the distribution still looks binomial (rather like a bell curve), centered about 8.14 or so.

If the Bone Man can measure to the nearest 1/100th, we should venture a guess of 8.14. That gives us the best chance of matching our guess with the Bone Man's actual average ratio.

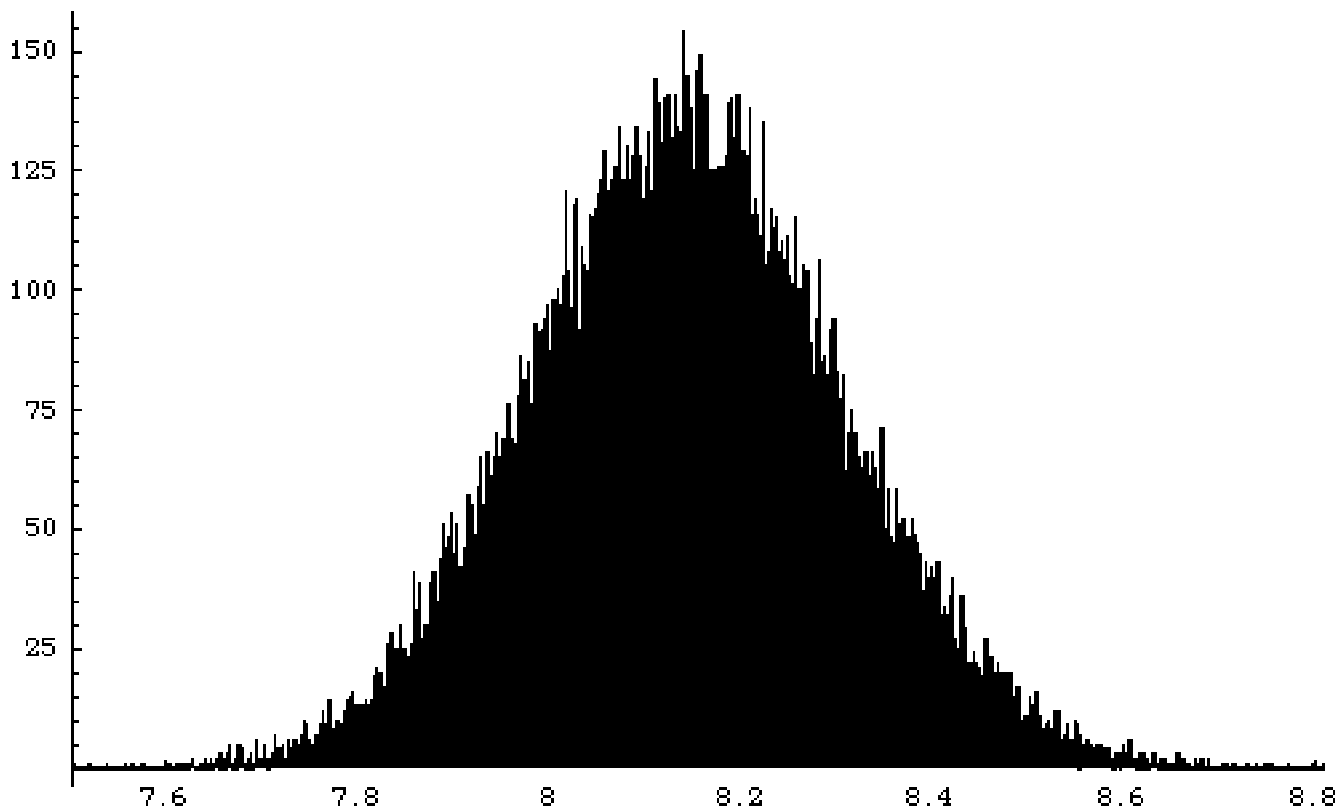

Distribution of Mean Ratios for measurements of 1/100th – 50,000 trials, bin size ~0.0013

1/1000th and 1/10000th

For completeness, I have also included below histograms of what happens if we could look at 50,000 different boneyards but with measurements of 1/1000ths and 1/10000ths. Note that according to the formula above, we SHOULD get that the most likely average ratios as 12.609 and 17.192 respectively. This is precisely where we see the peaks of the following distributions. Hooray, it looks like our predictions are working!

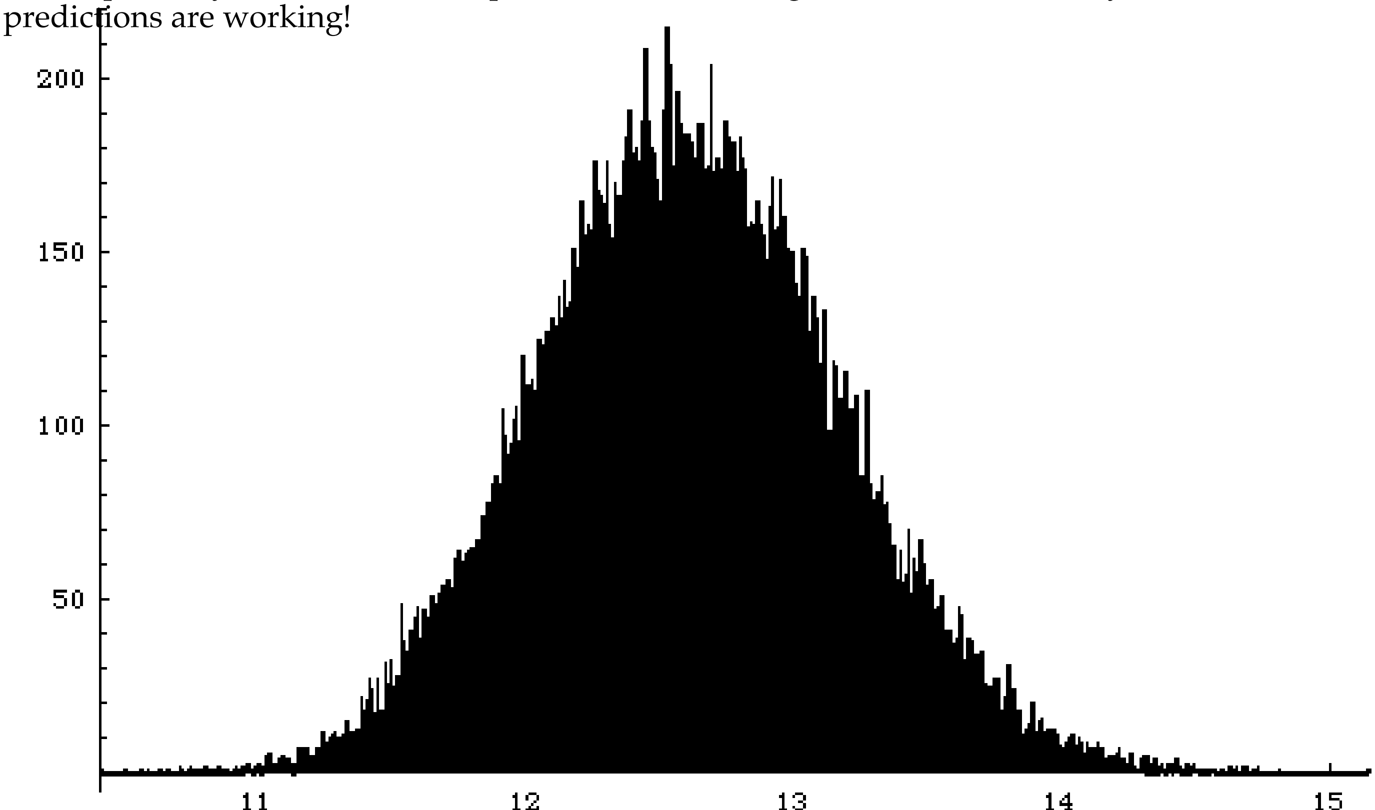

Distribution of Mean Ratios for measurements of 1/1000th – 50,000 trials, bin size ~0.005

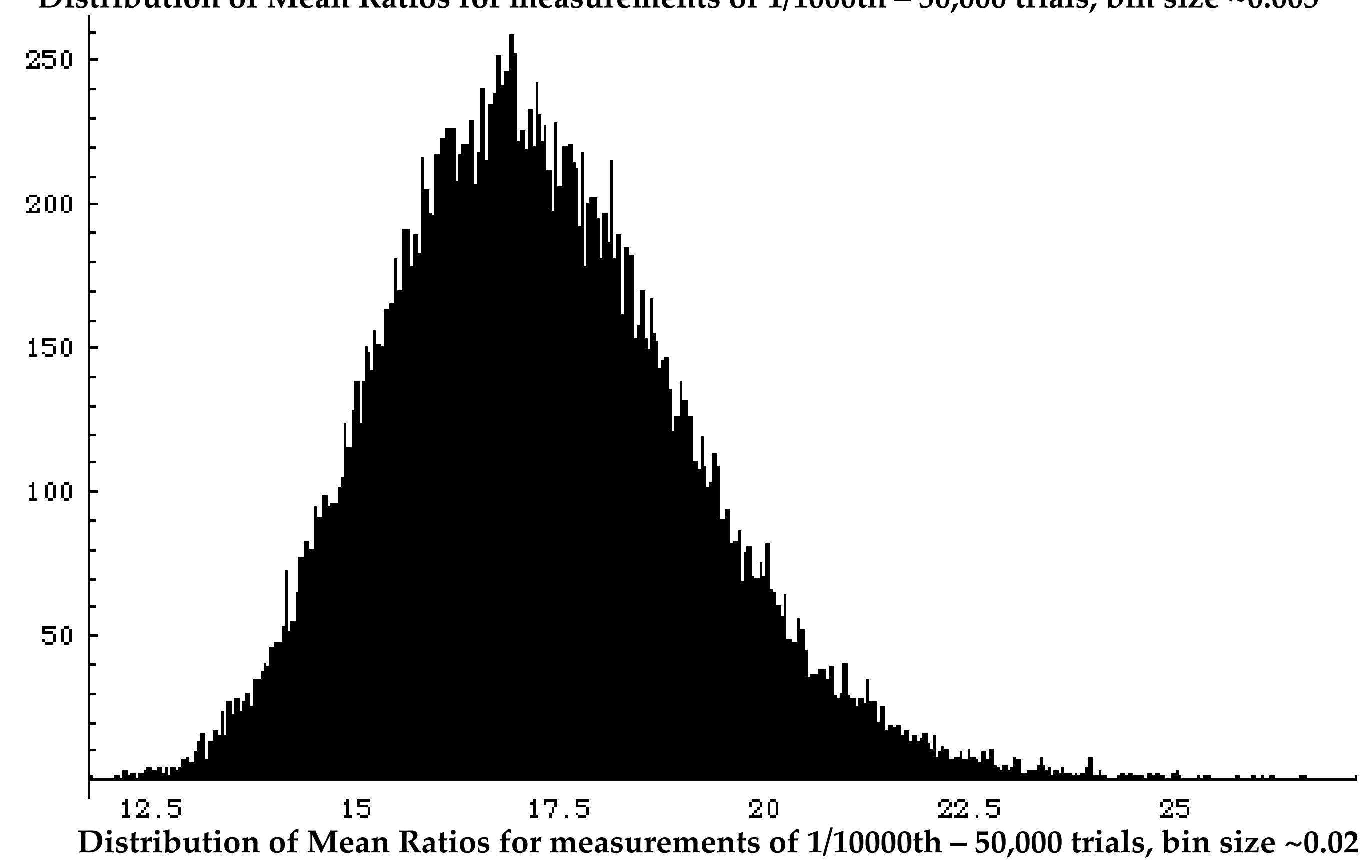

Distribution of Mean Ratios for measurements of 1/10000th – 50,000 trials, bin size ~0.02

Not So Fast

Let's take a look at the distribution we get for measurements in increments of 1/100,000ths. Actually, I'm getting tired of writing all of those zeroes. Let's call it "100thousandths". According to our formula, the most likely mean ratio should be 21.794. In fact, even in the simulation, when I ask Mathematica to calculate the mean of all of those simulated mean ratios, I get 21.788, which is basically right on. BUT – look at where the peak of the distribution is!

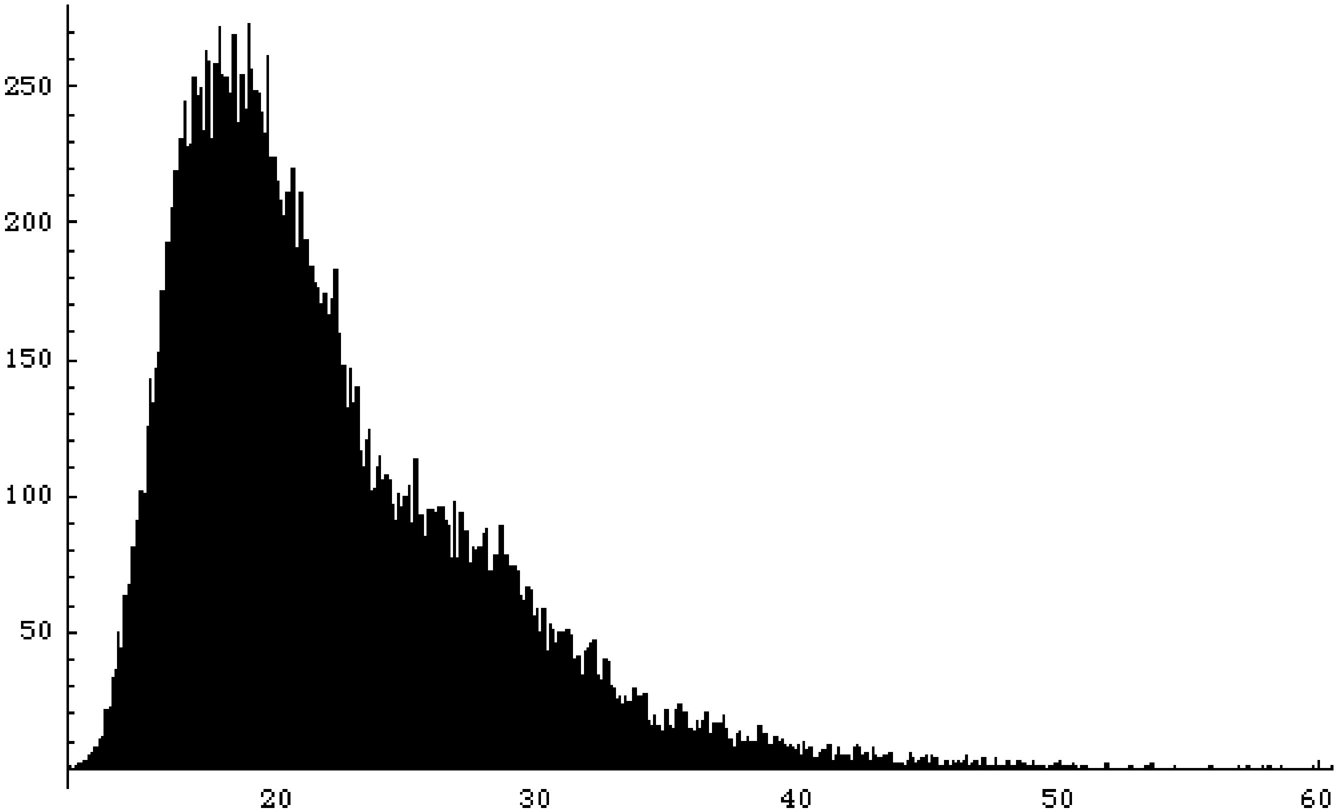

Distribution of Mean Ratios for measurements of 100thousandths – 50,000 trials, bin size ~0.05

The peak of the distribution is CLEARLY less than 20 and therefore the mean is not a good guess for us to make! But how can this be?

Unfortunately, we must revisit the places we started – multinomial distributions.

Let's think about the MOST LIKELY average ratio for measuring to 1/100th and having 10,000 bones. Just to find the most likely combination of bones breaks, we must evenly distribute the ratios among the 49 different bones break possibilities (.51, .52, .53, ..., .99). The multinomial coefficient for those 49 equally-probable different objects is found by

$$\frac{10{,}000!}{k_{51}!\cdot k_{52}!\cdot k_{53}!\cdots k_{99}!}\cdot\left(\frac{1}{49}\right)^{10{,}000}$$

where all of those k's added together must total 10,000 (from the number of bones). The way to make that fraction the largest possible is to make all of those k's as equal as possible. Think about it this way: if there were 49 bone break positions (as there are for measuring to the nearest 1/100th) for the 10,000 bones, then I am claiming that the largest possible multinomial coefficient occurs when forty-five of the bone break positions will be occupied by 204 bones and four of the bone break positions will be occupied by 205 bones. If the pile of bones were just slightly less evenly distributed, then either one of those 204-bone partitions would be reduced by one to then have only 203 bones in it and either one of the 205-bone partitions or one of the 204-bone partitions would be increased by one. This changes that part of the denominator from

204! * 205! to 203! * 206!

or

204! * 204! to 203! * 205!

Either way, the denominator increases. In the first case, we divide by 204 but we multiply by 206. In the second case we divide by 204 but multiply by 205. In either case, we have increased the denominator and so the overall coefficient decreases. So the most likely occurrence is for approximately equal numbers of bones to be as evenly distributed among the various bone break positions. That's the reason why my Monte Carlo simulations happen to be centered at the same place as the average ratio for measurements as fine as 10thousandths.

If we can measure to the closest 1/1000th, for instance, the most even distribution of bone ratios of those 10,000 bones will occur when 479 of the ratios are each found 20 times and twenty of the ratios are found 21 times.

If we can measure to the closest 10thousandths, for instance, the most even distribution of the bone ratios of those 10,000 bones will occur when 4997 different ratios are each found twice and two of the ratios are each found three times.

But what happens when we have too many possible ratios – many more possible ratios than we have bones? Well, the largest multinomial coefficients STILL occur with the most even distributions – but for 100thousandths measurements, that still leaves 10,000 different bone ratios each found once and 39999 ratios that are not found at all!

THIS it's the key to arriving at an actual answer to the problem: because the non-even distribution of the ratios (the bone breaks are evenly distributed instead), we are much less likely to see large ratios.

Let's see why this is. Examine the case of being able to measure to the nearest 1/10th. In this case, the largest ratio possible is 9 (0.9/0.1). The probability that that this ratio is present at all in the pile of 10,000 bones is:

$$P(9) = 1 - P(no9s) = 1 - \left(\frac{3}{4}\right)^{10,000}$$

$$P(9) = .99999999....$$

The probability has approximately 1250 nines in it – so it's pretty certain that there will be at least one ratio of 9 present in the pile of bones.

If our precision is increased to 1/100th, the maximum ratio we could have in the pile is 99 (0.99/0.01). The probability that we have a ratio of AT LEAST 9 present in the pile essentially does not change

$$P(\geq 9) = 1 - \left(\frac{39}{49}\right)^{10,000} = 0.999....$$

The probability that our maximum ratio possible is found in the pile is:

$$P(99) = 1 - \left(\frac{48}{49}\right)^{10,000} = 0.999...$$, with "only" approximately 90 nines.

So having a bone break ratio of 99 present in the pile is very probable. This large ratio, of course, shifts the average ratio upward, overshadowing most of the other bones in the pile.

For a precision of 1/1000th, the maximum ratio possible is 999. The probability that a ratio of 9 or greater or 99 or greater are each still essentially one. The probability that at least one 999 is present in the pile is:

$$P(999) = 1 - \left(\frac{498}{499}\right)^{10,000} = 0.999999998...$$

So it's still almost certain that at least one bone break in that pile of 10,000 bones will corresponding to the ratio 999.

For a precision of 1/10,000th, the maximum ratio possible is 9999. The probability that that this ratio is included in the pile of bones is:

$$P(9999) = 1 - \left(\frac{4998}{4999}\right)^{10,000} = 0.8647...$$

So now we see that while it is quite likely that a ratio of 9999 is present in the pile, it is by no means a certainty. If we continue on to see what the probability of having a ratio of 9999 being in the pile with more precise measurements, the probability does not change. It is always approximately 0.8647....

BUT, as we can see, the maximum ratio for a given accuracy is becoming less probable with increasing precision. When we get up to 1/15,000ths, the probability of a maximum ratio occurring is about 0.73. For 1/20,000ths, the probability of the maximum ratio being present in the pile is only 0.63.

So why does the peak of the distribution stop being centered on the mean of the distribution? Because the maximum ratios are no longer commonly found in the piles. As long as the number of bones is larger than one over the precision, the maximum ratio is essentially guaranteed of being found in the pile at least once. Because that maximum ratio is so large compared to any of the other ratios possible, the peak of the distribution must be fairly high; and, in fact, the peak of the distribution is centered on the mean of the distribution.

But when the number of bones is less than one over the precision, it becomes increasingly unlikely that the maximum ratio will occur – the peak would continue to move to higher ratios if those maximum ratios were sampled with any sort of regularity, but they aren't!

One may argue that the lowest ratios are hardly ever sampled as well; this is true, but no one misses ANOTHER ratio near 1, whereas missing a ratio near 9999999 makes sure that the distribution's peak stays where it is. Once the peak stops increasing, it makes sense that it stays where it is, because the probabilities of getting any single ratio or higher does not change with an increase in the precision. Since the probability of a ratio of 9999 occurring is 0.8647 for an accuracy of 1/10,000th, the probability is ALWAYS approximately 0.8647. Once we stop adding the higher ratios, since they are so unlikely, the peak of the distribution stops moving.

The one question remaining, then, is at point does the peak of the distribution stop moving?

To see what average ratio we should guess, I've zoomed in on the appropriate region of the histogram. It indicates that we should guess that the Bone Man's average ratio is 18.0. You will find these distributions on the following page.

Because the distribution is skewed, it is a bit difficult to name an exact peak. However, after doing some analysis on the zooms (varying the height at which I call the distribution symmetric and finding the center this way for many heights), I have found that all of the "centers" that I found occur between 17.9 and 18.1. Additionally, equation 1 is 18.0 when the precision is 1/15,000th. I find this to be more than coincidental, that the peak should stop when one over the precision is 1.5 times the number of bones. So using equation 1 and setting M=15,000, we find the peak should be at 18.007, so that is where I will call the center off the peak of the distribution.

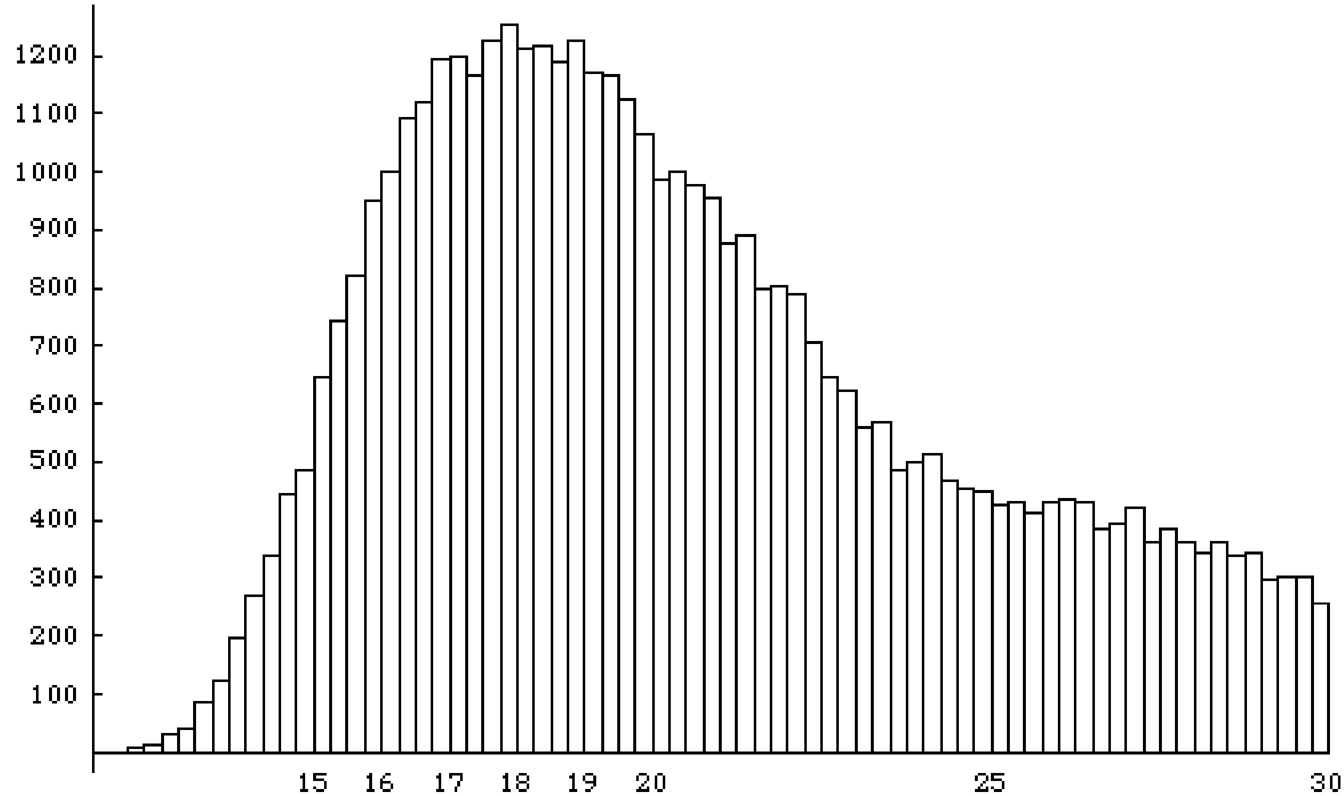

Distribution of Mean Ratios for measurements of 100thousandths – 50,000 trials, bin size 0.25

Distribution of Mean Ratios for measurements of 100thousandths – 50,000 trials, bin size 0.1

Even Larger Numbers
What follows are distributions from even more precise measurements (up to 100millionths).

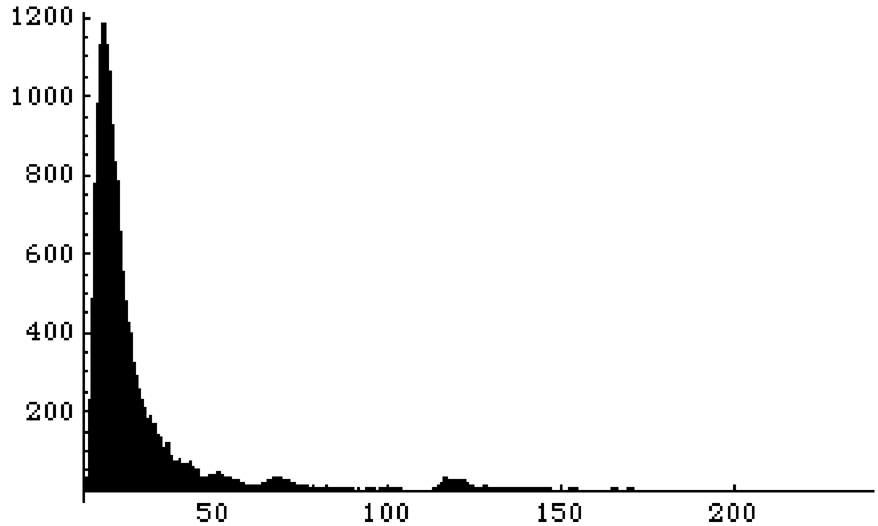

Distribution of Mean Ratios for measurements of millionths – 50,000 trials, bin size 0.1

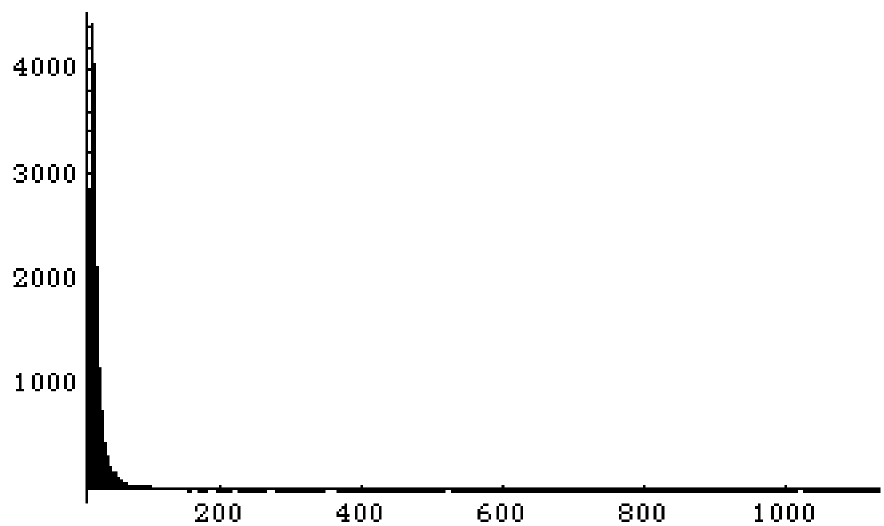

Distribution of Mean Ratios for measurements of 10millionths – 50,000 trials, bin size 0.1

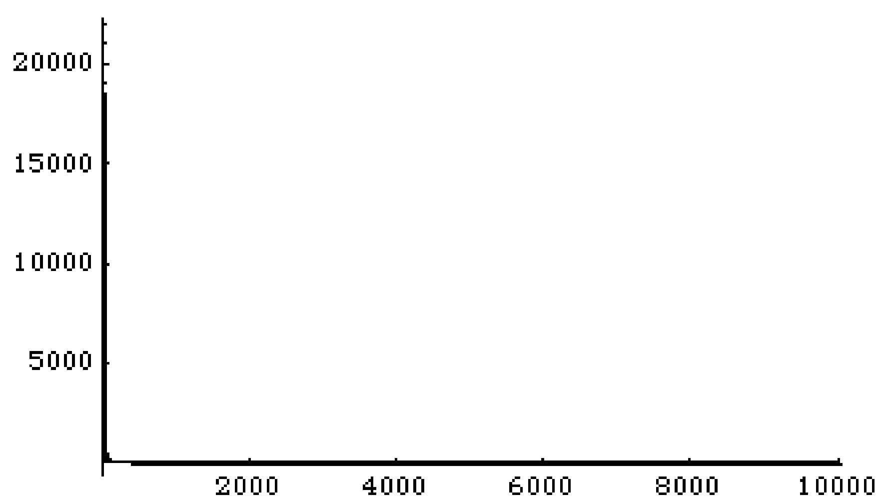

Distribution of Mean Ratios for measurements of 100millionths – 50,000 trials, bin size 0.1

Note how squished this last one is. Some average ratios has high as 10,000 occur, but most of them are much closer to the zero side of the histogram. Now I'll zoom in on each of these histograms to see where THEY have peaks.

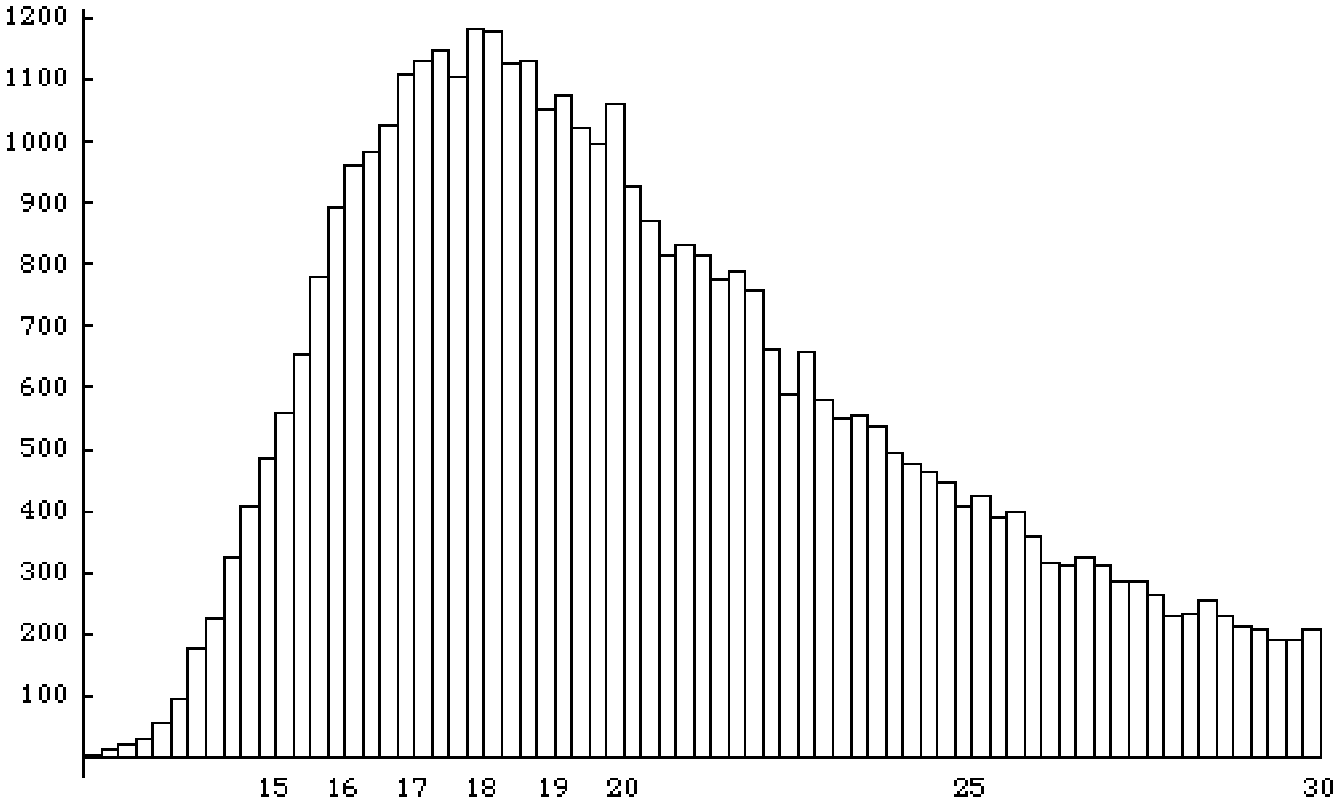

Distribution of Mean Ratios for measurements of millionths – 50,000 trials, bin size 0.25

Distribution of Mean Ratios for measurements of millionths – 50,000 trials, bin size 0.1

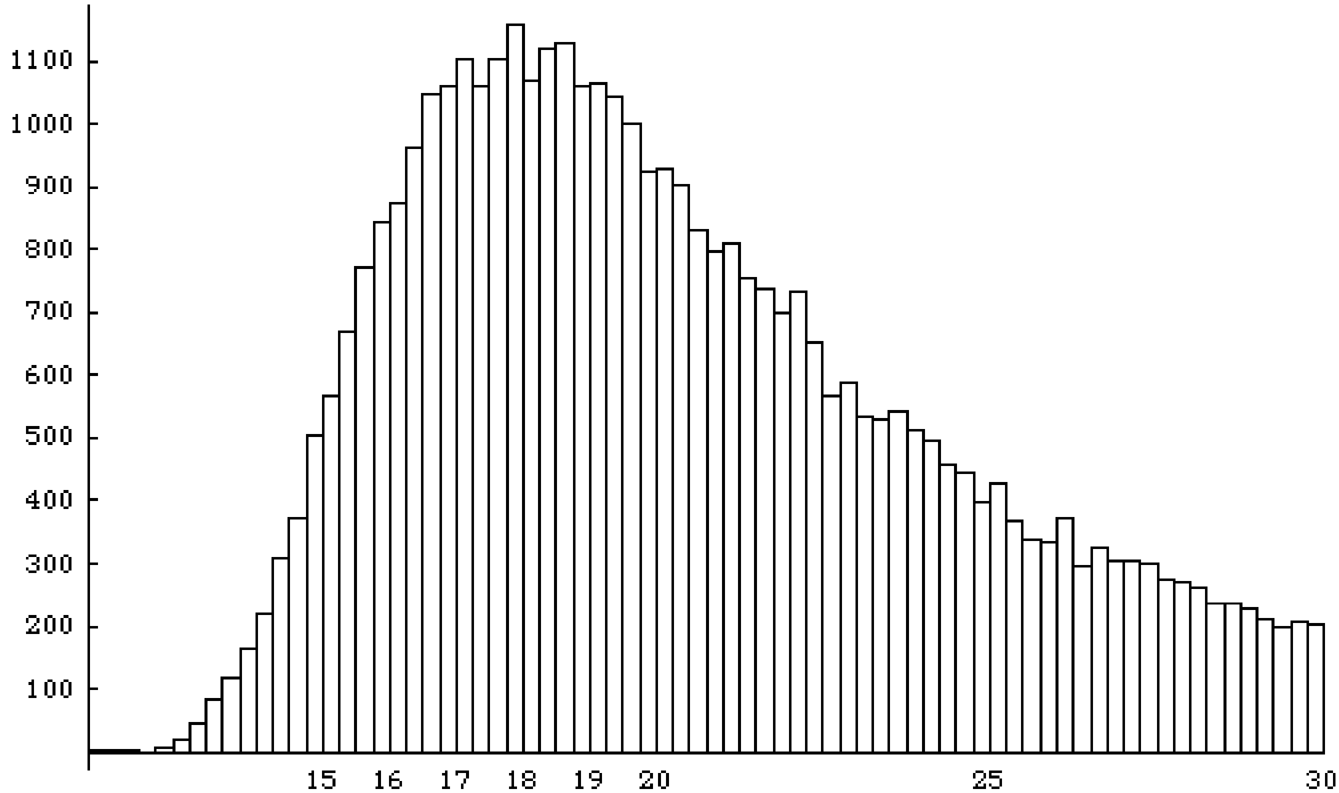

Distribution of Mean Ratios for measurements of 10millionths – 50,000 trials, bin size 0.25

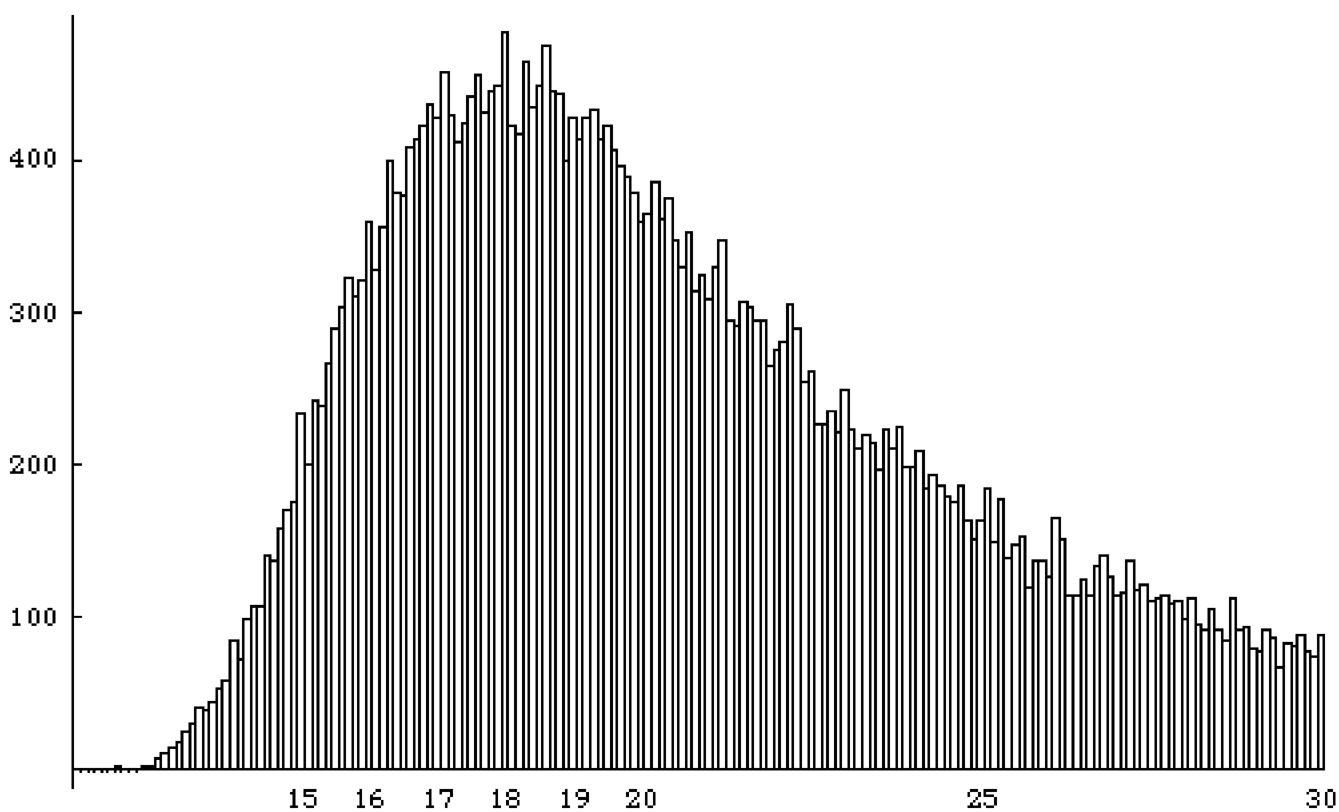

Distribution of Mean Ratios for measurements of 10millionths – 50,000 trials, bin size 0.1

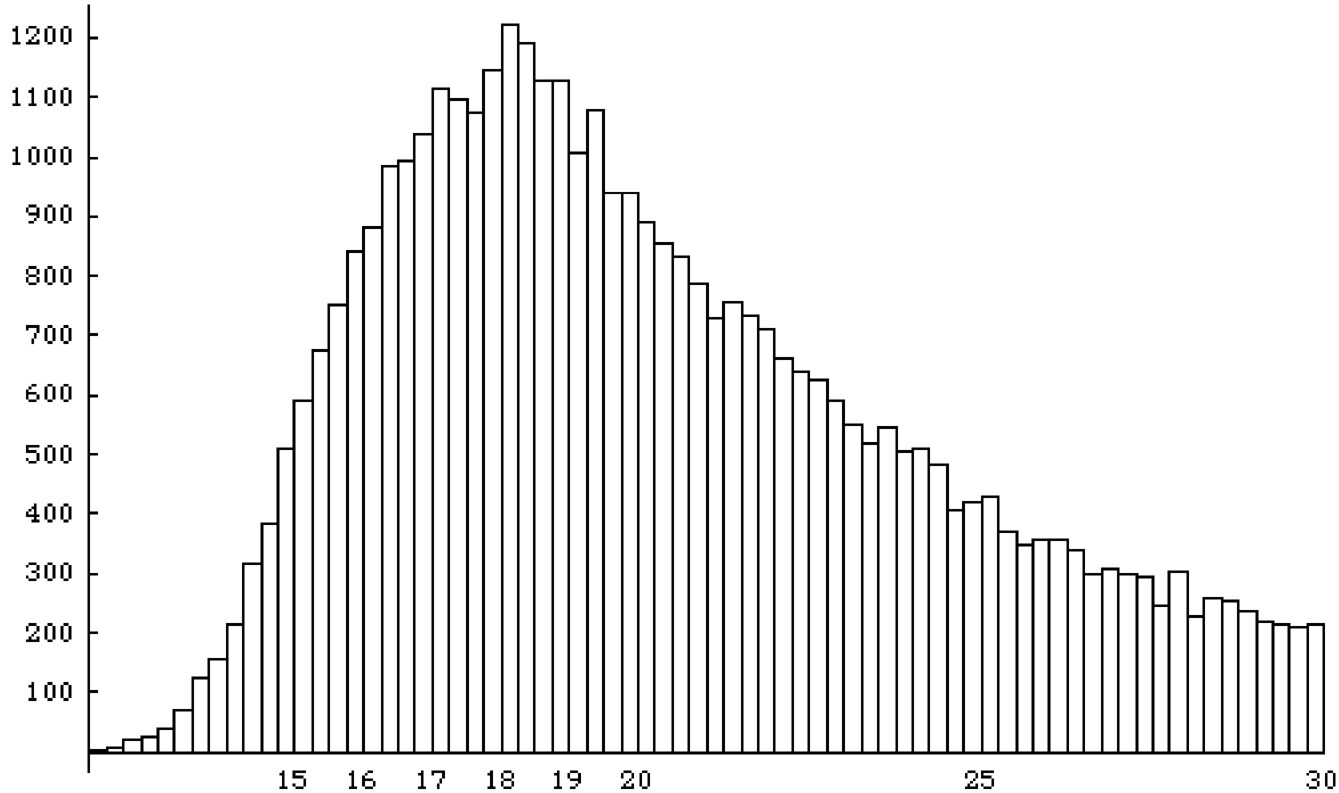

Distribution of Mean Ratios for measurements of 10millionths – 50,000 trials, bin size 0.25

Distribution of Mean Ratios for measurements of 100millionths – 50,000 trials, bin size 0.1

My suggestion, to anyone who has made it this far into the explanation, is to print those last three pages out and hold any two of them (one on top of the other) up to a light to compare the distributions. They are all nearly identical! In fact, in this "peak region," they are also almost all the same as the 100thousandths histogram (print that one out and check it, too). And they all center near 18.0.

The Number of Bones

To wrap up, I suppose I should explain WHY the distributions get stuck at 18.0. Unfortunately, I can't explain why it stops at EXACTLY 18.0. I have explained that the peak stops because for a given accuracy, the highest ratios do not often contribute to the mean ratio of the pile, but I do not have a magic numeric reason for why the pile stops moving EXACTLY when one over the precision is 1.5 times the number of bones.

To see if this type of trend holds, I looked to see what happens with piles of 100 bones. The peak stopped moving when M was approximately 150. When the piles had 100,000 bones, the peak stopped moving when M was approximately 150,000 bones. Eureka!

Conclusion

Simply put, you MUST ask the Bone Man how precise his measurement tool is. If it measures to a precision LESS precise than 1 in 10,000 units, use the formula (equation 1):

$$MostLikely = \frac{2 - M + 2\cdot M\cdot\gamma + 2\cdot M\cdot\Psi\left(\frac{M}{2}\right)}{M - 2}$$

If you don't have a calculator on hand that can deal with the Digamma function (and who does?), use this simple chart and make your best approximation:

Precision	Your Best Guess to save your leg
1/10th	4.208
1/100th	8.141
1/1000th	12.609
1/10,000ths	17.192
1/15,000ths	18.002
Anything More Precise	18.002

Let me finish by saying why the average ratios of 10,000 bones measured with infinite precision is NOT undefined. We can simply measure each of the bones and find the answer! The answer is only undefined if we have an infinite number of bones as well as an infinite number of breaking points.

I WILL say that although others have claimed (including myself!) that the answer to the question "What should you tell the Bone Man" is undefined (a much more reasonable assertion), I will now claim, given all the above work, that even with infinite precision you should guess 18.002. It appears, from my simulations, that the distributions near the peak do not change with an increase in precision (even over a wide range, like from 100thousandths to 100millionths). I wouldn't, therefore, expect that the distribution near the peak would change with infinite precision either. On the other hand, no matter what number you guess, if you have to be exact, your odds of keeping your leg are essentially nil. They're just a bit bigger than otherwise if you guess 18.002. If you can be within 0.5 of the ratio, then you're in pretty good shape.

The final formula for a measurement tool with precision 1/M, a number of bones B, and a "final" bone number $B_f = 3/2\,B$:

$$MostLikely(M) = \begin{cases} \dfrac{2 - M + 2\cdot M\cdot\gamma + 2\cdot M\cdot\Psi\left(\dfrac{M}{2}\right)}{M-2}, & M \le B_f \\[2ex] \dfrac{2 - B_f + 2\cdot B_f\cdot\gamma + 2\cdot B_f\cdot\Psi\left(\dfrac{B_f}{2}\right)}{B_f-2}, & M > B_f \end{cases}$$

To MY satisfaction, I'm done.

The answer is 18.002.