Michael Reiff

Reason and Rationality

September 10, 2025

The Cube Factory Paradox

"A precision tool factory produces iron cubes with edge length ≤ 2 cm. What is the probability that a cube has length ≤ 1 cm given that it was produced by that factory?" (van Fraassen 370).

This problem first appeared, in the formulation presented here, in the Dutch-American philosopher Bas van Fraassen's 1989 book *Laws and Symmetry*. It was, however, based on an earlier paradox proposed by the French mathematician Joseph Bertrand in his 1888 work *Calcul Des Probabilites.* Van Fraassen, in *Laws and Symmetry*, summarized Bertrand's paradox and his proposed solution this way: "the problem of choosing a number at random from [0, 100] is the same as that of choosing its square (*Calcul des probabilités*, 2nd edn., 4). [Bertrand] adds that these contradictions can be multiplied to infinity. His own conclusion is that when the sample space is infinite, the notion of choosing at random 'n'est pas une indication suffisante'—presumably not sufficient to create a well-posed problem" (van Fraassen 303). Soon you will see how the two paradoxes are related. But for now, suffice it to say that I believe Bertrand's conclusion is correct, and that the "solutions" to to both paradoxes are the same. My hope is that this essay can provide a clear, accessible explanation of why that is.

Before I do that, though, I must answer the question of why this matters. Bertrand's paradox (and therefore also van Fraassen's) were created to demonstrate that the "principle of indifference"—the view that in a scenario where you have no evidence regarding the probabilities of the various possible events, you should assign an equal probability to each possibility—cannot always be correct. This has wide ranging implications, including for Bayesian probability, which involves establishing "priors," or probabilities before a certain piece evidence is considered. With that, let's turn to the paradox itself.

To visualize the problem, you can imagine having two boxes, one for cubes with side lengths ≤ 1cm and the other for cubes with side lengths greater than 1cm. If the factory produces randomly sized cubes, it seems obvious that there is a 50% chance that any given cube will go in the first box, and a 50% chance it will go in the second.

But what exactly is meant by "randomly sized?" Surely if the cubes have random *side lengths*, my original analysis will be correct. But what if the cubes actually have random *volumes*?

In this case, the side lengths would no longer be random (and therefore trending toward an even distribution between the two boxes), but would skew toward 2 cm. The reason for this can be seen most clearly with a graph.

Graph of y = x^(1/3):

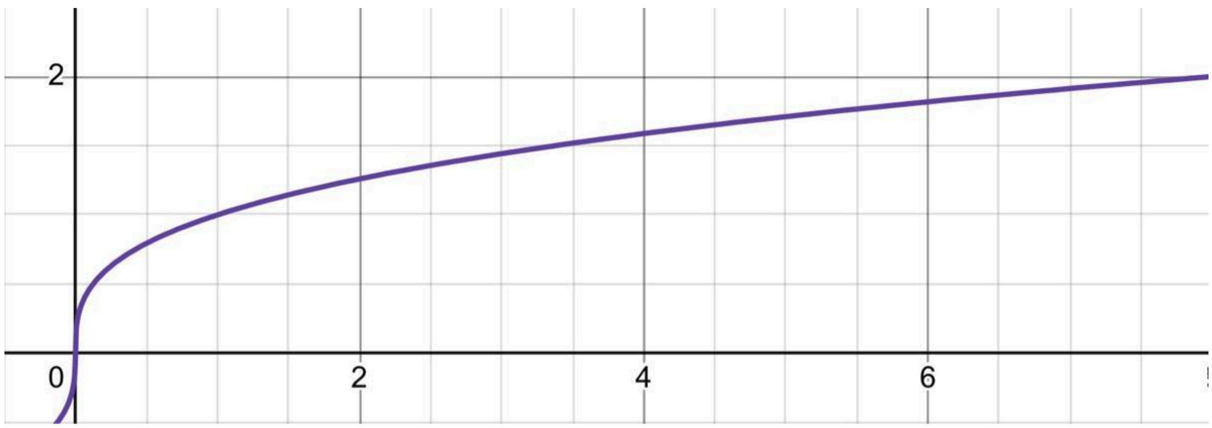

x = volume (cubic cm)

y = side length (cm)

Looking at the graph, which shows the side length (y) associated with any volume (x) between 1 cubic cm and 8 cubic cm, it's obvious that there are more values of x for which their corresponding y-values will be between 1 and 2 than between 0 and 1 (see below)*.

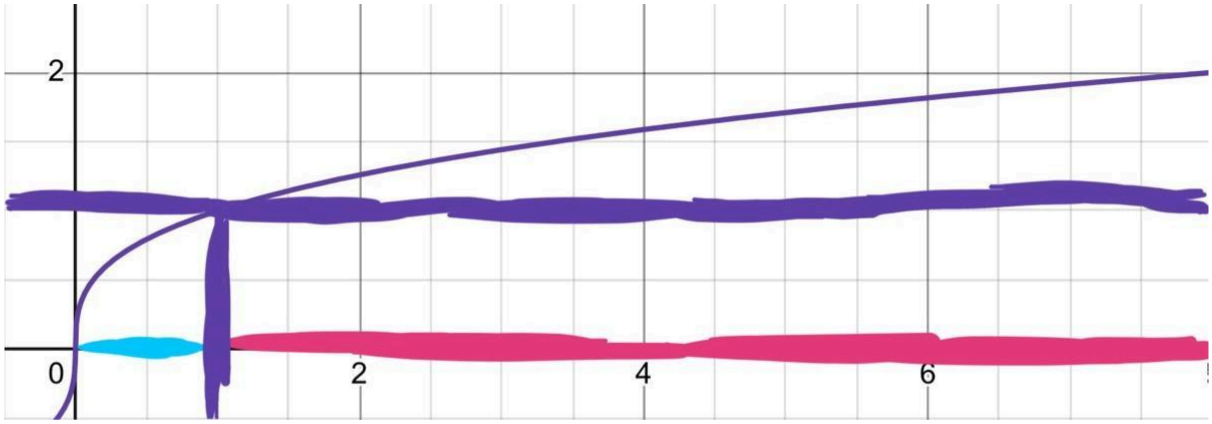

Indeed, “randomly sized” could also mean that the faces of each cube have random areas, which would result in yet another unique distribution of side lengths.

Graph of y = x^(1/2):

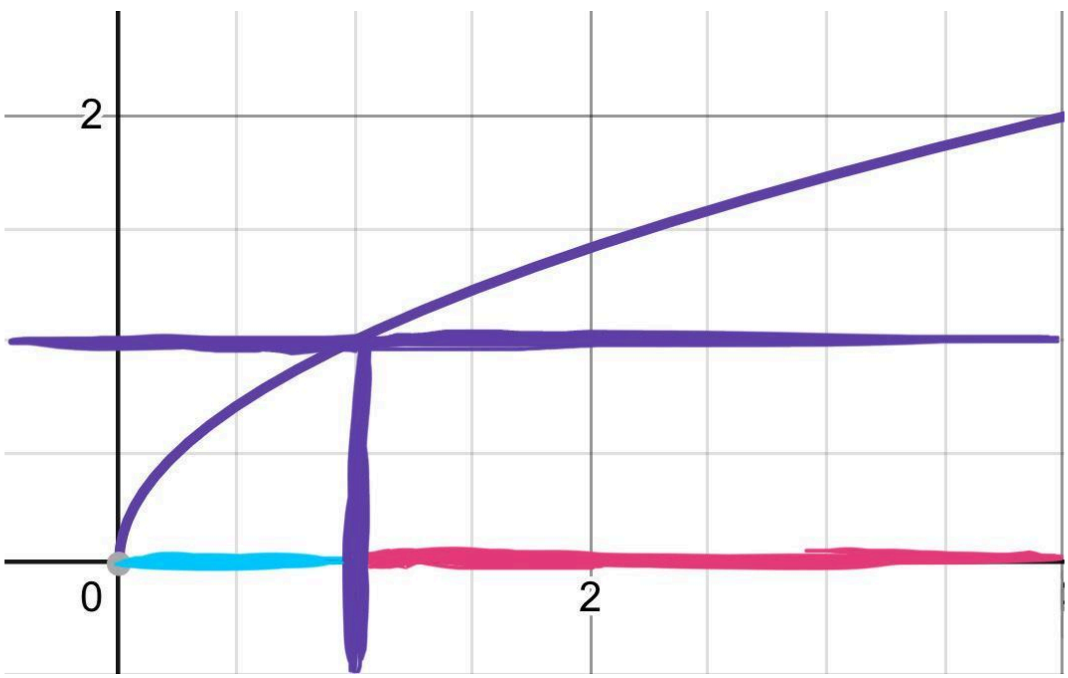

x = area (square cm)

y = side length (cm)

Just as each cube has a 2-dimensional area and a 1-dimensional length associated with it (that is, an area or length which could become the cube using only exponentiation), each cube could also have a 4-dimensional tesseract associated with it. Thus, you could also imagine interpreting “randomly sized” to mean having a random 4-dimensional “hypervolume.”

Graph of y = x^(1/4):

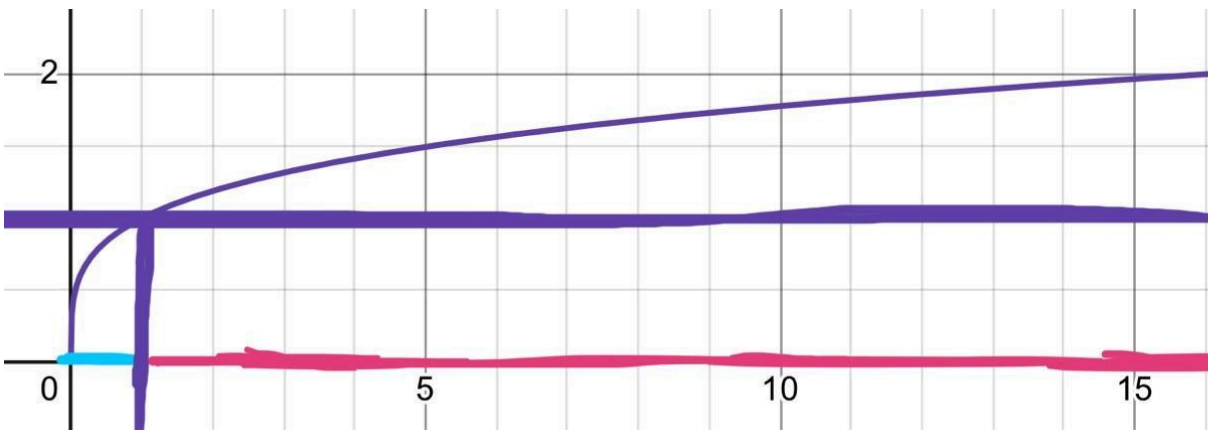

x = 4-dimensional hypervolume (quartic cm)

y = side length (cm)

Thus, there are an infinite number of ways “random size” could be legitimately interpreted, as you could continue to increase the dimensions of the hypervolume that you’re assigning to be random indefinitely. It also seems like each time this is done, a different blue:red ratio will result. That means there’ll be a different ratio of side lengths ≤ 1 cm to side lengths > 1 cm.

In fact, this must be true, because the equation for these graphs in general is y = x^(1/n), where n is the dimension of the unit being used in the x-axis.

The blue line always starts at 0 and continues along the x-axis until it reaches the x value of the coordinate where the graphs of y = x^(1/n) and y = 1 intersect (i.e., the point where the side lengths become > 1). Therefore, the x value at which the blue line stops is the solution to the equation 1 = x^(1/n) for any given value of n. Raising both sides of this equation to the y power results in the equation 1^y = x. Since 1 raised to any power equals 1, x will always equal 1, regardless of what n equals. Thus, the blue line will always stop when it reaches 1. Additionally, because the blue line always starts at 0, the blue line will always have a length of 1 (whichever until is being used in the x-axis).

The red line always starts where the blue line stops, and continues along the x-axis until it reaches the x value of the coordinate where the graphs of y = x^(1/n) and y = 2 intersect (i.e. the point where the side lengths become greater than 2). Therefore, the x value at which the red line stops is the solution to the equation 2 = x^(1/n) for any given value of n. If you graph this equation, it's clear that, as n increases, x will also increase.

Note: This is actually the graph of 2^n = x, because Desmos was having a hard time with 2 = x^(1/n). I don't think it effects the argument at all, but just remember that technically this graph is slightly incorrect, because n can never = 0. The way to prove that x and n are positively correlated in the original equation without the graph is to simply notice that, whenever n increases, the exponent as a whole (which is its reciprocal) will cause x to decrease. Thus, in order to keep the expression equal to 2, x must increase.

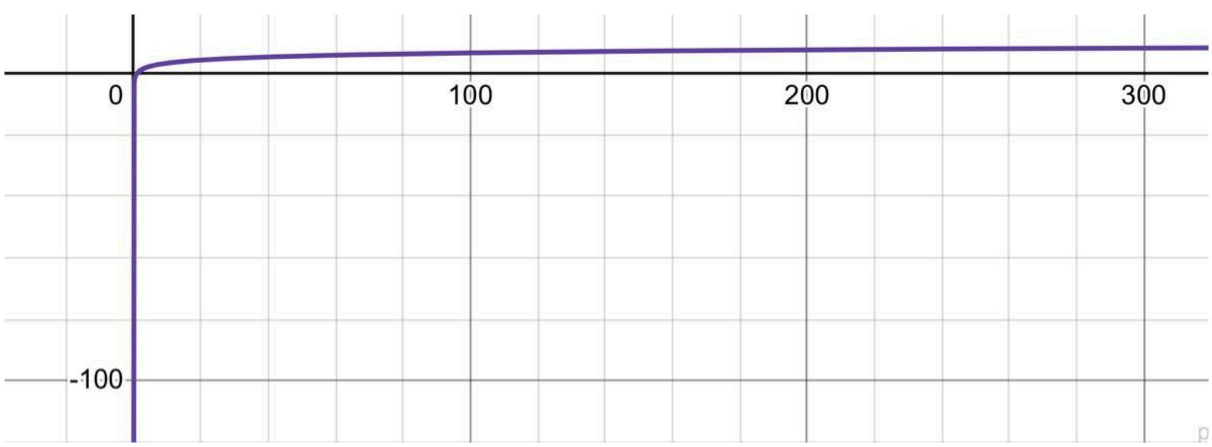

x = the value at which the red line stops

n = the dimension of the unit used in the x axis of the preceding graphs

Since the length of the red line increases as n increases and the length of the blue line is always 1, it must be true that, as n increases, the ratio of the red line's length to the blue line's length increases. Therefore, the ratio of values of x for which their corresponding y-values will be between 1 and 2 to values of x for which their corresponding y-values will be between between 0 and 1 also increases. Since the values of x are random (and therefore trend toward an even distribution), such an increase would result in an increase in the ratio of side lengths > 1 cm to side lengths ≤ 1 cm toward which the values of y (that is, side lengths) would trend.

Now, if every time n increases, this ratio also increases, then no two "n"s could result in the same ratio. Therefore, since there are an infinite number of "n"s (i.e., dimensions that can be used in the x-axis—or, more fundamentally, different legitimate interpretations of "random size"), there must also be an infinite number of ratios of cubes in the first box to cubes in the second box toward which the cubes could trend.

So far, I've shown that there are an infinite number of possible legitimate interpretations of "random size," each of which result in a different ratio of cubes between the two boxes toward which the cubes would trend. To give a definitive probability that any given cube would fit into the first box, therefore, you would first have to decide on the probability of each possible interpretation of the question being correct. But if you try to do that, you'll run into some trouble.

To see why, start by imagining a simpler case: A random number is about to be generated. What's the probability that it will be a 5, for example? Already there's some confusion. Would 5.1, for instance, "count as" 5? If the answer is no, then 5 must really mean 5.0. But then you could just ask the same question with 5.11, and so on. So eventually you'll need to conclude that 5 in this case really means 5.000… with an infinite number of zeros following the 5. If you say that the answer is yes, however, then it seems like some randomly generated numbers are more likely than others, and thus a truly random number is not being generated. But you're trying to examine a situation where it is, so you must take the "infinite zeros" approach. This also leads to problems, though. If each number had an equal non-zero probability of being generated, than adding all of the probabilities up would result in infinity. But this is impossible, because the probabilities, when added together, must equal 1. So the probability of any given number being generated must be 0. But it seems like this doesn't work either, because adding all of these probabilities up results in 0, meaning the total probability of any number being generated is 0. It therefore seems impossible to generate a random number in this sense**.

Likewise, it now seems impossible that you could generate a random n value, and therefore, impossible to generate a cube of a truly "random size." Nevertheless, it's surly possible for a factory to produce a cube with an edge length ≤ 2 cm. So you must conclude that just because you have no evidence of what size the next cube will be, it doesn't necessarily mean that every possibility is equally probable.

*One Possible Objection:

"There are an infinite number of numbers on the blue line and an infinite number of numbers on the red line. You claim that a random number from 0 to 8 being on the red line is more probable, but since there are already an infinite number of numbers on each line, it seems like the red line being longer wouldn't really change anything."

My Response:

The "probability density" (the probability / the length of the range) for the range 0 to 8 is 1/8, because 8 (the length of the range) * 1/8 (the probability density) = 1 (the probability, which *must* = 1 because x *must* be between 0 and 8). The probability density will remain 1/8 for any smaller range of x-values because, if the probability density were lower for some ranges and higher for others, the probability of a random x-value falling within a particular range would vary for different ranges of the same length, which can't be true because no x value is more probable than any other, and probability equals the length of the range * the probability density.

Since the length of the range * the probability density = the probability, the probability of the randomly generated x-value being in the blue section is 1 * 1/8, or just 1/8, and the probability of it being in the red section is 1 * 7/8, or just 7/8.

Remember that on the graph, the probability of the number that gets picked being within a given range is represented by area, not length. Insofar as the blue and red lines are acting as widths of rectangles, they are certainly different lengths, which result in different areas despite each being multiplied by the same hight value (namely, 1/8).

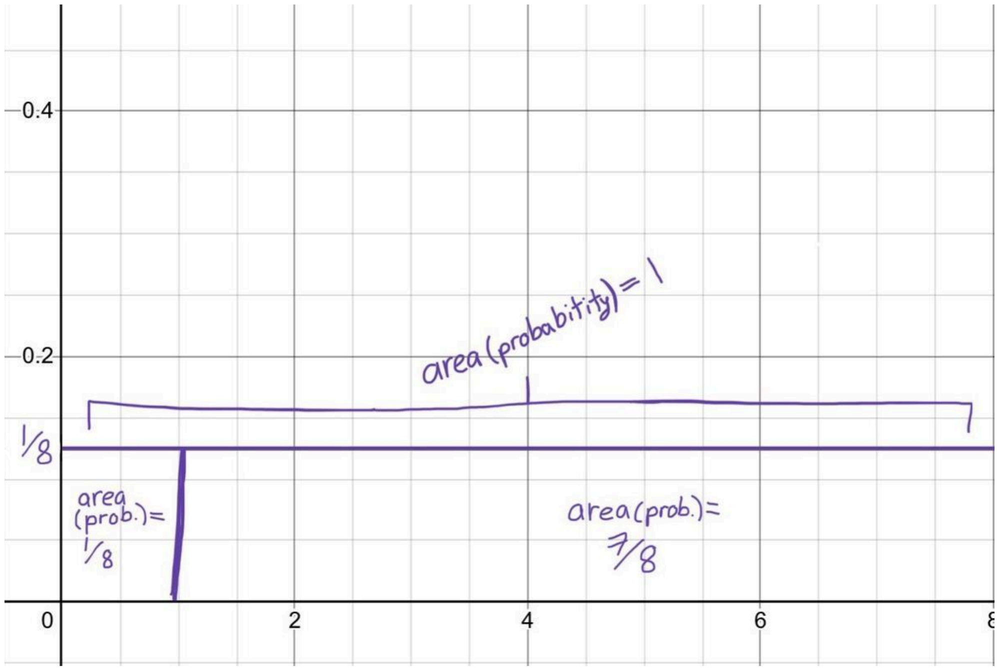

**If the random number were being generated within a set range (say, from 1 to 10), you would at least be able to say that each number has a non-zero "probability density", and that this kind of random generation is therefore possible, even if there's an infinite amount of numbers within that range. But since there is no set range of possible values, and probability density is equal to the probability divided by the length of the range, calculating the probability density in this case would be impossible. In other words, infinite possibilities is not the problem here, but rather a lack of a defined range of acceptable values.

Works Cited

Bertrand, Joseph. *Calcul Des Probabilités*. 1888.

van Fraassen, Bas C. *Laws and Symmetry*. Clarendon Press, 1989.