# The infinite antinomy

## Common types of number

There are seven common types of number that most of us learn about in school:

**1. Natural numbers:** The natural numbers are the counting numbers (0, 1, 2, 3, 4, 5, …). They answer questions like "how many fingers am I holding up?" In 1889, the distinguished Italian mathematician, Giuseppe Peano, published a formal definition of the natural numbers that is widely accepted today. Basically they are defined as a sequence, starting at zero, that is generated by adding 1 to the previous number in the sequence.

**2. Integers:** The integers include the natural numbers together with the sequence of numbers, starting at zero, in which each number is 1 *less* than the one before.

**3. Rational numbers:** The rational numbers include the natural numbers and integers, with the addition of the fractions that result from dividing one integer by another. It may be worth noting that dividing any rational number by another always amounts to dividing one integer by another. For example:

$$1\frac{2}{7} \div 2\frac{5}{6} = \frac{54}{119}$$

Expressing fractions in base-10 decimal expansion often results in a never-ending sequence of digits. For example, 1/3 = 0.33333333333… (the digit 3 keeps repeating indefinitely). However, a decimal fraction expansion of a rational number always contains some sequence of digits that repeats regularly, which distinguishes it from an irrational number (the same is also true if you use any other base than 10).

**4. Irrational numbers:** An irrational number is a non-integer whose fractional expression requires a never-ending sequence of digits that contains no repeating pattern.

**5. Real numbers:** The real numbers are the rational and irrational numbers together. Real numbers that can't be defined algebraically are known as *transcendental* numbers.

**6. Imaginary numbers:** An imaginary number is a real number multiplied by the imaginary unit $\boldsymbol{i}$, which is equal to the square root of $-1$.

**7. Complex numbers:** A complex number is a real number added to an imaginary number.

# All the natural numbers

Take the natural numbers— all of them. They're clearly defined (each being one larger than the one before), so let's collect them in a complete set that includes them all.

That sounds reasonable, except how do you get them all? The natural numbers go on indefinitely, without an upper bound. There can't be a largest natural number, because if there were one, you could create a bigger one by adding 1 to it. They are countless, because the largest one, if it existed, would be their count. So, when you think about it, we don't know, and indeed can't know, how to take all of them. No matter how many you take, there's still more. We haven't taken all of them until we get to the end, and they're endless.

The solution is to assign them an upper bound, namely *infinity* (∞), much in the same way we assign the imaginary unit `i` to be the number that when multiplied by itself equals -1. However, let's be clear: *infinity is not itself a number; it's a contradiction in terms*. It's the end of a sequence that has no end, the stopping point of a process that has no stopping point. It attempts to encompass the entirety of something that cannot be encompassed.

Nonetheless, the concept of infinity is useful in practice for completing all kinds of asymptotic convergences. When you can show that the more you repeat a process, the closer you get to a result without actually reaching it, you can reach that result by invoking infinity. Infinity allows you to complete what can't be completed in any finite number of steps.

We use infinity this way all the time. In the rational numbers, for instance, we commonly equate a precise fractional value to an endless repeating sequence of decimal digits. For example, we accept that `0.33333...` converges on and ultimately equals `1/3`, or that `0.99999...` converges on and ultimately equals `1`. We're invoking infinity in very much the same way when we say, "the set of all natural numbers."

Although you could define the natural numbers declaratively, as an infinite sequence starting at zero, every element of which is one greater than the preceding one, you'd still be including the concept of infinity in the definition ("infinite sequence"). To keep infinity separate, Peano defined the natural numbers imperatively, as generated by a process that is repeated over and over. This seems to avoid the unknowable somewhat better, unless you want to generate all the natural numbers.

Think what it would take to compute a large number of them. Even if you set things up so that you could add one to the preceding number in a trillionth of a second, you'd still use up all the time in the universe before you reached infinity. And, long before that, you would have reached such an unimaginably vast number that all the data centers currently in existence wouldn't have the capacity to store it.

Generating an infinite sequence is simply impossible in practice.

It's important to repeat this: infinity can't be reached. It's fundamentally unknowable, and intrinsically contradictory. There's no reasonable way to think about it. If you want to add two infinities together, do you start the second one at the end of the first?

# So let's get real

Georg Cantor, father of set theory, tried to keep infinity's ghastly inconsistency under control. It was important, logically, to be able to take the infinite set of all natural numbers without including infinity. Rather than getting stuck trying to reach the end of an endless process, he worked with the idea of combining two endless processes into a single process that merges them. For example, he noted that there are ways to put every possible rational fraction in a one-to-one correspondence with every possible natural number. An example of such a mapping begins:

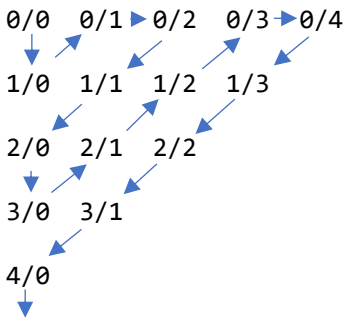


The first 12 correspondences in this sequence are as follows:

| | | | | | |
|---|---|---|---|---|---|
| 0 | → | 0/0 | 8 | → | 1/2 |
| 1 | → | 1/0 | 9 | → | 0/3 |
| 2 | → | 0/1 | 10 | → | 0/4 |
| 3 | → | 0/2 | 11 | → | 1/3 |
| 4 | → | 1/1 | 12 | → | 2/2 |
| 5 | → | 2/0 | 11 | → | 3/1 |
| 6 | → | 3/0 | 12 | → | 4/0 |
| 7 | → | 2/1 | and so on... | | |

Interestingly, fractions like 1/0, 2/0, and 3/0 are not generally included in the rational numbers because they are infinite (any fraction that has a non-zero numerator approaches infinity as the denominator approaches zero). We generally sidestep infinity in that case by saying more politely that the result of dividing anything by zero is indeterminate.

Because you can "*count*" the number of rational numbers in ways such as the one illustrated above, Cantor believed that the infinity of rational numbers is the same "size" as the infinity of natural numbers. This logic is almost universally accepted today. A contrary school-child might disagree, because the natural numbers in the mapping above increase much, much faster than the numerators and denominators of the fractions they're counting, suggesting that there are more rational numbers than natural numbers. Of course, that's only true until you get to the end of the process, which by definition you can't.

Taking this kind of *countability* as an indicator of infinite equivalence, Cantor used his famous "diagonal argument" to prove that you can't count real numbers the same way. From this he concluded that the infinity of real numbers is **larger** than the infinity of natural numbers.

# Cantor's diagonal argument

The diagonal argument goes as follows: With the real numbers between zero and one expressed as binary (base-2) fractional expansions, suppose you set up a one-to-one correspondence *of any kind* between them and the natural numbers. Your list of real numbers would look something like this:

| Naturals | | Real numbers |
|---|---|---|
| 0 | → | 0.**0**10101010101... |
| 1 | → | 0.0**1**1010101010... |
| 2 | → | 0.11**0**011001100... |
| 3 | → | 0.001**0**01001001... |
| 4 | → | 0.0001**1**1000111... |
| 5 | → | 0.10110**1**011010... |
| ... | | . . . . . . |

The diagonal argument points out that no matter how you've made your list of real numbers, you can always find a real number that's not on it by taking the $\mathtt{i}^{th}$ digit of the $\mathtt{i}^{th}$ number in the list for all `i` and then changing `0`s to `1` and `1`s to `0` in the number that that diagonal sequence specifies. That resulting number then differs from every single number on your list by at least one digit.

This is obviously true for any finite list, but *is it also true for an infinite list*? Do we know what happens when the list keeps going indefinitely? You can observe in a finite list of this sort that the farther you go down it, the better finite approximations of your "new" diagonal number you come to, although you can never reach that new number at any finite stopping point because you're always creating a new and different number there. But what does that actually mean?

For Cantor, it meant that the infinity of real numbers is larger than the infinity of integers. This reasoning is built into the foundation of modern mathematics and mathematical logic. More generally, Cantor extended the idea to the set of subsets of any infinite set. Finite sets of size `n` have $2^n$ subsets, and presumably infinite sets have $2^\infty$ subsets too. Cantor argued that the infinity that measures the size of the set of subsets of any infinite set is a larger infinity than the one that measures the size of the set itself, and this conclusion let him posit the existence of an infinite hierarchy of so-called transfinite numbers.

One outcome of this set-theoretic argumentation is to let you take the set of all natural numbers without including infinity in it. I assert, however, that the inconsistency inherent in the concept of infinity actually makes this impossible. In the same way we agree that the infinite sequence 0.99999… is equal to 1, we should be able to agree that an infinite set of distinct numbers must contain at least one infinite number (and if it contains one, it must contain an infinite number of them).

But how about that diagonal argument? Let's consider a simple and obvious mapping of natural numbers to real numbers in the unit interval [0,1). This is the part of the real number line that starts at 0 and goes up to, *but does not include*, 1.

To create the mapping, take the binary digits that express each natural number and reverse them on the other side of the decimal point to express real binary fractions in that interval, as follows:

| Base10 | | Binary | | Real fractions |
|---|---|---|---|---|
| 0 | = | 0 | → | 0.000000000000... |
| 1 | = | 1 | → | 0.100000000000... |
| 2 | = | 10 | → | 0.010000000000... |
| 3 | = | 11 | → | 0.110000000000... |
| 4 | = | 100 | → | 0.001000000000... |
| 5 | = | 101 | → | 0.101000000000... |
| 6 | = | 110 | → | 0.011000000000... |
| 7 | = | 111 | → | 0.111000000000... |
| 8 | = | 1000 | → | 0.000100000000... |
| 9 | = | 1001 | → | 0.100100000000... |
| 10 | = | 1010 | → | 0.010100000000... |
| 11 | = | 1011 | → | 0.110100000000... |
| ... | | ... | | . . . |
| ∞ | = | ...111111 | → | 0.111111111111... = 1 |

If you continue this mapping infinitely, the number that the diagonal argument says won't be on your list after you change all those 0s to 1s is 0.11111111....

Is that number on your list? Well, it starts out as a legitimate fraction less than one, but it converges on 1, which we've explicitly excluded from the interval.

If you "*complete*" the mapping above, you reach both the infinite sequence of digits that isn't supposed to be on your list and also the limiting number that isn't supposed to be in your interval. Plus, of course, you've also created *every possible infinite sequence of binary digits* along the way.

This suggests two conclusions that Cantor's logic (and mathematical logic today) manages to avoid:

1. The finite concept of an open interval doesn't properly apply to the real number line. You can't take a complete bounded interval of real numbers without including its bounding points just as you can't take the set of all natural numbers without including infinity.

2. Trying to distinguish different sizes of infinity doesn't work: Infinity can't be confined like that. The infinite sequence of natural numbers is capable of containing or encompassing any conceivable number including infinite numbers of infinite infinities raised to the powers of an infinite tower of infinities. We have to stop trying to make sense of that.

The effort to treat an infinite set like a finite one and get around infinity's contradictory nature has created problems in set-theoretic logic for a long time. Bertrand Russell and Alfred North Whitehead made a heroic effort early in the twentieth century to build a comprehensive logical foundation for mathematics based on set theory, but they gave up. Their efforts foundered on a set-theoretic antinomy, a paradox like Epimenides' *this statement is false* (if it's true, it's false, but if it's false, it's true).

Russell illustrated his set-theoretic antinomy like this: If there's a set of all people who don't shave themselves, and there's a barber who shaves everyone in that set but no one else, *is the barber in that set*? If he is, he isn't, but if he isn't, he is. Gödel's incompleteness theorems and Turing's computability extension of them are related to this kind of underlying set-theoretic antinomy.

I'd like to suggest that infinity is Russell's barber.

# Irreal numbers

If you set aside Cantor's transfinite numbers, it's interesting to consider the infinite set of infinite natural numbers that are expressed by different infinite sequences of digits.

I propose calling these numbers the *irreals*.

Thinking about them, you can see why Cantor might have wanted to exclude them. On the one hand, irreal numbers are clearly distinct from one another. They are not all the same number. Each is represented uniquely by a *different* infinite sequence of digits.

On the other hand, they don't really have a "size." Any one of them can encompass any other, or all the others. Finite differences between them can't affect their relative size in this sense, since any finite number divided by infinity converges to zero.

Paradoxically, however, they can be compared to one another in a way that suggests ordinality.

Consider this one, for example, expressed in decimal digits:

$$\sum_{n=1}^{\infty} 10^n = \ldots 11111111110$$

Although this infinite sum doesn't converge on any value we're used to talking about, you can see that we can still add 1 to it several times to produce a sequence in the same way we do when generating other sequences of natural numbers:

```
...11111111110
...11111111111
...11111111112
...11111111113
```

In fact, irreal numbers can be understood as an ordered continuation and completion of the natural numbers. If we include irreals in the generation of rational numbers, the distinction between rationals and reals goes away.

Now consider a second irreal, defined by multiplying the first one by three:

$$\sum_{n=1}^{\infty}(3 \times 10^{n}) = \ldots 33333333330$$

This irreal appears to be greater than the other one ($> \ldots 11111111110$), even if it doesn't differ in "size". For any *finite* value of n, the second sum is clearly larger than the first. Also, if you subtract $\ldots 11111111110$ from $\ldots 33333333330$, you get:

$$\sum_{n=1}^{\infty}(2 \times 10^{n}) = \ldots 22222222220$$

At the same time, finite ordinality doesn't apply to these numbers, because infinity's indeterminate nature intrudes. If a finite limit of n in the first sum were to exceed a finite limit of n in the second, even by 1, the first number would be greater than the second. If infinity is the limit of n in both cases, we can't say that their limits are equal, because infinity encompasses every number that's bigger than you can specify.

Because we have to acknowledge that infinity is not a single determinate value, the size of irreal numbers is also indeterminate, even if they seem to have some ordinality. To repeat: their size is simply limitless.

# Ghost numbers

We can only access irreal numbers whose infinite sequence of digits we have some way of defining. There are infinitely many that must exist, but except where there's a pattern or algorithm that precisely defines them, we can never know what they are. This, of course, is also true for real numbers. As Gregory Chaitin has pointed out, infinitely many of the real numbers in any finite interval on the number line can *only* be expressed as an arbitrary infinite sequence of digits that no formula or algorithm can define. Such numbers (and the points on the line that they represent) are *completely inaccessible* to us. We know they must be there, but it's absolutely impossible to identify one. I call them *ghost numbers*.

Ghost numbers make up the infinite set of real and irreal numbers that can only be defined by indeterminable infinite sequences of digits, and they constitute a permanent hole in our ability to know. They aren't constrained as to the distribution or mean value of their digits, so the law of large numbers does not apply. Some of them can be partially compressible, as long as an infinite part of their expression is indeterminable. For example, the decimal expression of a ghost fraction might start with 0.5, followed by ten billion zeros, followed by an indeterminable infinite sequence of digits. The number is still a ghost number, with an unknowable value that happens to be quite close to one half.

# Real intervals

Real (and irreal) numbers define the real number line, which is one-dimensional and has unit length as its measure. For any two real numbers a and b where a is less than b, the interval on the real number line from a to b has these characteristics:

- The interval contains an infinite number of points, no matter how close a is to b.
- It is *complete* in that it includes *every* real number greater than a and less than b.
- The *length* of the interval is equal to (b – a).

If you make the interval incomplete by removing even the smallest line segment from it, then the combined lengths of the two intervals that remain won't add up to (b – a) any more.

Suppose you remove a single point, with zero length, from the interval, does that still leave an overall length equal to (b – a)? Well, yes, but as I've already pointed out, you can't really exclude any single point, because its neighbors converge on it infinitely. This is a problem that Cantor tried to get around by letting the infinite set of finite natural numbers be defined without including infinite numbers.

What you can do, though, is to remove an infinite number of line segments in such a way as to extract their bounding points, as the next section describes.

# The Triadic Cantor set

Cantor was intrigued by what is now generally called the triadic Cantor set, first described by Henry John Stephen Smith, an Irish mathematician.

It's good to use base-3 numbers to describe it, for reasons that will become clear. It's composed of points on the real number line in the unit interval between `0` and `1`.

Specifically, it's the set of points defined by this process:

1. Remove the middle third of the interval between `0` and `1`, which is the segment from `0.1` up to `0.2` in base 3 numbering,.
2. Remove the middle third of the remaining two segments (from `0.01` up to `0.02` and from `0.21` up to `0.22` in base-3 numbering).

Here's what the results of these 2 steps look like:

If you continue to repeat this process of removing the middle third of remaining segments *infinitely*, the triadic Cantor set is what's left when you're done.

Expressing the segments in base-3 lets you define what we're doing declaratively: we're taking away every single point along the interval whose base-3 representation contains the digit **1** anywhere in it. These points define an infinite number of line segments.

We can add up the total length, L, of those segments, like this:

$$L = \frac{1}{3} + \frac{2}{9} + \frac{4}{27} + \frac{8}{81} + \cdots$$

The same thing can be expressed like this:

$$L = \sum_{n=0}^{\infty} \frac{2^n}{3^{n+1}}$$

If you multiply both sides of these expression by two thirds, you get:

$$\tfrac{2}{3}L = \tfrac{2}{9} + \tfrac{4}{27} + \tfrac{8}{81} + \cdots \quad \textit{or} \qquad \tfrac{2}{3}L = \textstyle\sum_{n=1}^{\infty} \frac{2^n}{3^{n+1}}$$

You can see now that two thirds of $L$ is the same infinite sequence of lengths that follows 1/3 in the first equation. If you subtract two thirds of $L$ from both sides of these equations, you get:

$$\tfrac{1}{3}L = \tfrac{1}{3} \qquad \text{or:} \quad L = 1.$$

So, the combined lengths of all the segments we've removed add up to the entire length of the interval, namely **1**.

What's left? After taking away every number between **0** and **1** whose expression in base 3 contains a **1** anywhere in it, we're left with every number between **0** and **1** that can be expressed in base 3 using only **2**s and **0**s. This is the triadic Cantor set. However, if you were to replace the **2**s with **1**s in every number in the set, you'd have a complete binary (base-2) expression of every point in the interval we started with. We can map the triadic Cantor set to the entire infinite set of natural numbers expressed in binary notation like this:

```
      1   →   0.2
     10   →   0.02
     11   →   0.22
    100   →   0.002
    101   →   0.202
    111   →   0.222
   1000   →   0.0002
    ...   →     ...
...11111   →   0.22222...
```

And of course, being infinite, the set also contains an infinite number of ghost numbers.

The triadic Cantor set is interesting because even though you can find points in it that are arbitrarily close together, it contains no intervals: its points are completely disconnected.

Between any two of them, no matter how close together they are, lie an infinite number of points that don't belong to the set.

It's also useful to consider that although the process of removing middle thirds of intervals appears to remove all points from the line except the Cantor set, it also appears to remove the entire line. This seems contradictory until you realize that the Cantor set consists entirely of the bounding points of the segments you remove, and as I argued above, you can't properly exclude bounding points from an interval. They're actually still there.

## Totally random

Okay, let's try something different: pick a random number in the range zero to fifteen.

What am I asking? It's not like requesting a prime number, because numbers don't have an inherent quality of randomness. Randomness comes into play only in how you choose them. It's the *choosing process* that's random.

The best way of defining the randomness of a process is that you can't predict its outcome based on any available information. Note that it's a practical, heuristic definition, not a deductive one.

A good coin flip does satisfy that definition. It generates a number between zero and one that you can't predict based on information available to you (unless you happen to be psychic). Using it four times, you can generate a binary representation of the random number between zero and fifteen I was requesting— and there's no information available that lets you predict what it will be in advance.

The next question is, if you generate a sequence of numbers using a random process like a coin flip, does the sequence itself have any intrinsic property of randomness?

The randomness of the generating process makes knowing part of the sequence no help in predicting any of its other elements. The trouble is, that's not an intrinsic characteristic that you can test. The best you can say is that you don't see a pattern in the sequence, but there might be a pattern you missed. Or, you might see a very clear pattern that was generated by chance, with no actual predictive power.

That fuzzy empirical definition of a sequence's randomness has been refined over time to be more rigorous (but still empirical) by restating it in terms of the complexity or information density of the sequence. It still comes down to an empirical definition that says a sequence is random to the extent it's incompressible— meaning that you can't find any way to define it more concisely than by listing out its elements (but there might be a way you've missed).

Still, it's the most useful definition we have for randomness in a sequence. It also allows sequences to be partially random. For example, in a sequence of groups of ten numbers, the first nine elements in each group might always be the same while the tenth was generated by a random process. The sequence is then obviously compressible, and many of its elements are completely predictable, but part of it is likely hard to compress. A real-world example of partial randomness would be Brownian motion, where a particle jiggles around in a random walk— you

can never predict exactly what its next location will be, but each new location depends heavily on where the particle was before.

A fundamental problem here is that a random process like a coin flip is perfectly capable, totally by chance, of generating a finite sequence having an easily discernable pattern.

# Infinitely lucky

The stability of our world rests on probability distributions, right down to the quantum level. Entropy evens everything out, keeping the air pressure in a room pretty uniform. We rely on things mostly turning out the way they're likely to. Except, of course, the great universal river of entropy has back-eddies like life, like human civilization.

The thing is, when it comes to very large systems and/or very long time-spans, the likeliness of unlikely events increases. A popular example is the idea that an army of monkeys might produce the works of Shakespeare, given enough time spent pecking randomly at typewriters. A less popular but no less bizarre example is the Boltzmann-brain thought experiment, suggested in 1896 by physicist Ludwig Boltzmann. He argued that it's as likely for a brain to form spontaneously out of random atoms, purely by chance, as it is for the universe to be organized the way we think it is. Yes, the formation of such a brain is extraordinarily unlikely, but it's theoretically *possible*.

Take that to the limit and the contradictory nature of infinity can make anything that's possible certain, not to mention what's downright impossible.

Take a simple example: let's imagine using a random process like coin flipping to generate a sequence of zeros and ones. Probability predicts that the resulting sequence is more and more likely to have the same number of zeros as ones the longer we keep going. It's true that in the first ten flips we could get seven zeros and three ones, but over ten thousand flips we're very likely indeed to get about five thousand zeros and five thousand ones.

Keep in mind, though, that although chance favors that uniform distribution over time, it doesn't *guarantee* it. The odds of getting seven zeros in a row, for example, is one in $2^7$, or one in 128. Flipping a coin ten times provides three opportunities to get seven zeros in a row (the first three flips), so the odds become three in 128. Flipping 131 times gives you even odds, and flipping 259 times gives you two to one odds in favor of getting those seven zeros in a row.

Now let's imagine that we can run this process an infinite number of times, generating an infinite sequence of zeros and ones. Mathematicians imagine this kind of thing all the time. As the mathematician Vladimir Arnold is said to have observed, math is just like physics, except the experiments are so much cheaper.

The odds of getting 30 zeros in a row are one in $2^{30}$, or about one in a billion. The odds of getting a thousand zeros in a row are one in $2^{1000}$, which is so unlikely as to be practically impossible. However, when you run the coin-flipping process an infinite number of times, those odds become $\frac{\infty}{2^{1000}}$, which means that you're absolutely *certain* to generate that extraordinarily

unlikely sub-sequence in your infinite series of coin tosses. In fact, you'll get it not once, but an infinite number of times.

In fact, somewhere in your infinite sequence, any finite pattern of ones and zeros that is possible *must* exist. For example, say you interpret sub-sequences of 16 binary digits as Unicode alphabetic characters. Then, somewhere in your infinite sequence you'll be certain to find every single sequence of words that humans have ever written or ever will write, in every language that the Unicode glyphs can represent. You'd also find all the brilliant novels that will never be written, and scientific papers explaining all the greatest mysteries of the universe. And, of course, you'd find infinite quantities of utter garbage and meaningless gobbledygook. None of it would be accessible in practice, but the thought experiment once again suggests the power of infinity to subvert what we feel is reasonable.

Although any finite sub-sequence is bound to show up in your infinite sequence, the odds that your entire infinite sequence will consist only of zeros is $\frac{1}{2^{\infty}}$, or zero— it's just not possible.

But wait, we've only run our imaginary infinite coin-flip process once, generating a single infinite sequence. Let's go two-dimensional, and run the infinite process an infinite number of times. Now the odds of generating an infinite sequence of zeros totally by chance becomes $\infty/2^{\infty}$, which no longer looks like zero, but is also not calculable. Or is it 1?

## Well, so what?

Like it or not, infinity is deeply embedded in our concept of number. We can't get rid of it even if we want to. What we can do is fully acknowledge how self-contradictory it is, and how deeply it disrupts our ability to reason.

In physics, we're pretty comfortable intuitively with how things work at a local level. Only when we look at a much larger or smaller scale do things look different than sensible people expect them too (I've been told that despite all appearances the earth isn't flat).

The same may be true of mathematics, which would confirm a remark attributed to John von Neumann: we don't actually understand things, we just get used to them.