Ok, ok, we get it: what are logarithms about?

Logarithms find the cause for an effect, i.e. the input for some output

A common "effect" is seeing something grow, like going from $100 to $150 in 5 years. How did this happen? We're not sure, but the logarithm finds a possible cause: A continuous return of $\frac{ln\,(\frac{150}{100})}{5} = 8.1\%$ would account for that change. It might not be the actual cause (did all the growth happen in the final year?), but it's a smooth average we can compare to other changes.

By the way, the notion of "cause and effect" is nuanced. Why is 1000 bigger than 100?

- 100 is 10 which grew by itself for 2 time periods ($10 \cdot 10$)
- 1000 is 10 which grew by itself for 3 time periods ($10 \cdot 10 \cdot 10$)

We can think of numbers as outputs (1000 is "1000 outputs") and inputs ("How many times does 10 need to grow to make those outputs?"). So,

$$1000 \text{ outputs} > 100 \text{ outputs}$$

because

$$3 \text{ inputs} > 2 \text{ inputs}$$

Or in other words:

$$\log(1000) > \log(100)$$

Why is this useful?

Logarithms put numbers on a human-friendly scale.

Large numbers break our brains. Millions and trillions are "really big" even though a million seconds is 12 days and a trillion seconds is 30,000 years. It's the difference between an American vacation year and the entirety of human civilization.

The trick to overcoming "huge number blindness" is to write numbers in terms of "inputs" (i.e. their power base 10). This smaller scale (0 to 100) is much easier to grasp:

- power of 0 = 10^0 = 1 (single item)
- power of 1 = 10^1 = 10
- power of 3 = 10^3 = thousand
- power of 6 = 10^6 = million
- power of 9 = 10^9 = billion
- power of 12 = 10^{12} = trillion
- power of 23 = 10^{23} = number of molecules in a dozen grams of carbon
- power of 80 = 10^{80} = number of molecules in the universe

A 0 to 80 scale took us from a single item to the number of things in the universe. Not too shabby.

Logarithms count multiplication as steps

Logarithms describe changes in terms of multiplication: in the examples above, each step is 10x bigger. With the natural log, each step is "e" (2.71828...) times more.

When dealing with a series of multiplications, logarithms help "count" them, just like addition counts for us when effects are added.

Show me the math

Time for the meat: let's see where logarithms show up!

Six-figure salary or 2-digit expense

We're describing numbers in terms of their digits, i.e. how many powers of 10 they have (are they in the tens, hundreds, thousands, ten-thousands, etc.). Adding a digit means "multiplying by 10", i.e.

$$1[1 \text{ digit}] \cdot 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10[5 \text{ more digits}] = 10^5 = 100,000$$

Logarithms count the number of multiplications *added on*, so starting with 1 (a single digit) we add 5 more digits (10^5) and 100,000 get a 6-figure result. Talking about "6" instead of "One hundred thousand" is the essence of logarithms. It gives a rough sense of scale without jumping into details.

Bonus question: How would you describe 500,000? Saying "6 figure" is misleading because 6-figures often implies something closer to 100,000. Would "6.5 figure" work?

Not really. In our heads, 6.5 means "halfway" between 6 and 7 figures, but that's an adder's mindset. With logarithms a ".5" means halfway in terms of multiplication, i.e. the square root (9^.5 means the square root of 9 -- 3 is halfway in terms of multiplication because it's 1 to 3 and 3 to 9).

Order of magnitude

We geeks love this phrase. It means roughly "10x difference" but just sounds cooler than "1 digit larger".

In computers, where everything is counted with bits (1 or 0), each bit has a doubling effect (not 10x). So going from 8 to 16 bits is "8 orders of magnitude" or 2^8 = 256 times larger. (These bit sizes refers to the amount of memory available, not the processor speed). Going from 16 to 32 bits means 16 orders of magnitude, or 2^{16} ~ 65,536 times larger.

Isn't "16 extra bits of memory" better than "65,536 times more memory?"

Interest Rates

How do we figure out growth rates? A country doesn't intend to grow at 8.56% per year. You look at the GDP one year and the GDP the next, and take the logarithm to find the *implicit* growth rate.

My two favorite interpretations of the natural logarithm (ln(x)), i.e. the natural log of 1.5:

- Assuming 100% growth, how long do you need to grow to get to 1.5? (.405, less than half the time period)
- Assuming 1 unit of time, how fast do you need to grow to get to 1.5? (40.5% per year, continuously compounded)

Logarithms are how we figure out how fast we're growing.

Measurement Scale: Google PageRank

Google gives every page on the web a score (PageRank) which is a rough measure of authority / importance. This is a logarithmic scale, which in my head means "PageRank counts the number of digits in your score".

So, a site with page rank 2 ("2 digits") is 10x more popular than a PageRank 1 site. My site is PageRank 5 and CNN has PageRank 9, so there's a difference of 4 orders of magnitude (10^4 = 10,000).

Roughly speaking, I get about 7000 visits / day. Using my envelope math, I can guess CNN gets about 7000 * 10,000 = 70 million visits / day. (How'd I do that? In my head, I think 7k · 10k = 70 · k · k = 70 · M). They might have a few times more than that (100M, 200M) but probably not up to 700M.

Google conveys a lot of information with a very rough scale (1-10).

Measurement Scale: Richter, Decibel, etc.

Sigh. We're at the typical "logarithms in the real world" example: Richter scale and Decibel. The idea is to put events which can vary drastically (earthquakes) on a single 1 - 10 scale. Just like PageRank, each 1-point increase is a 10x improvement in power.

Decibels are similar, though it can be negative. Sounds can go from intensely quiet (pindrop) to extremely loud (airplane) and our brains can process it all. In reality, the sound of an airplane's engine is millions (billions, trillions) of times more powerful than a pindrop, and it's inconvenient to have a scale that goes from 1 to a gazillion. Logs keep everything on a reasonable scale.

Logarithmic Graphs

You'll often see items plotted on a "log scale". In my head, this means one side is counting "number of digits" or "number of multiplications", not the value itself. Again, this helps show wildly varying events on a single scale (going from 1 to 10, not 1 to billions).

Moore's law is a great example: we double the number of transistors every 18 months (image courtesy [Wikipedia](https://en.wikipedia.org)).

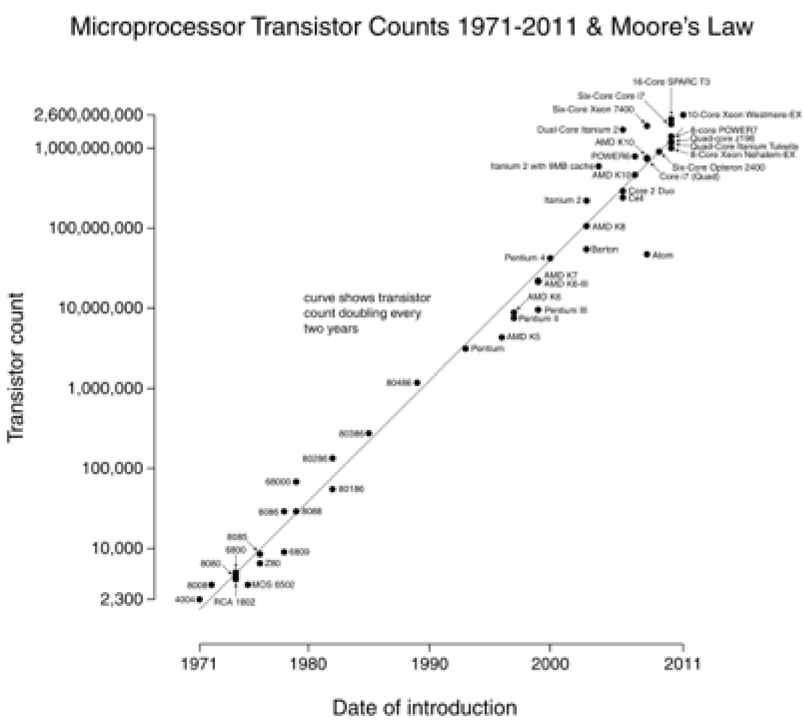

The neat thing about log-scale graphs is exponential changes (processor speed) appear as a straight line. Growing 10x per year means you're steadily marching up the "digits" scale.

Onward and upward

If a concept is well-known but not well-loved, it means we need to build our intuition. Find the analogies that work, and don't settle for the slop a textbook will trot out. In my head:

- Logarithms find the root cause for an effect (see growth, find interest rate)
- They help count multiplications or digits, with the bonus of partial counts (500k is a 6.7 digit number)

Happy math.

How to Think With Exponents and Logarithms

Here's a trick for thinking through problems involving exponents and logs. Just ask two questions:

Are we talking about inputs (cause of the change) or outputs (the actual change that happened?)

- Logarithms reveal the inputs that caused the growth
- Exponents find the final result of growth

Are we talking about the grower's perspective, or an observer's?

- e and the natural log are from the grower's instant-by-instant perspective
- Base 10, Base 2, etc. are measurements convenient for a human observer

In my head, I put the options in a table:

	Find Cause	Find Effect
From Grower's Viewpoint	$\ln(x)$ *Natural Logarithm*	e^x *Exponent base e*
From Observer's Viewpoint	$\log_{10}(x),\ \log_2(x)\ \dots$ *Another logarithm*	$10^x, 2^x \dots$ *Another exponent*

I have thoughts like *"I need the cause, from the grower's perspective… that's the natural log."* (Natural log is abbreviated with lowercase LN, from the high-falutin' *logarithmus naturalis*.)

I was frustrated with classes that described the inner part of the table, the raw functions, without the captions that explained when to use them!

That won't fly, let's get direct practice thinking with logs and exponents.

Scenario: Describing GDP Growth

Here's a typical example of growth:

- From 2000 to 2010, the US GDP changed from 9.9 trillion to 14.4 trillion

Ok, sure, those numbers show change happened. But we probably want insight into the cause: What average annual growth rate would account for this change?

Immediately, my brain thinks "logarithms" because we're working backwards from the growth to the rate that caused it. I start with a thought like this:

$$\text{logarithm of change} \rightarrow \text{cause of growth}$$

A good start, but let's sharpen it up.

First, which logarithm should we use?

By default, I pick the natural logarithm. Most events end up being in terms of the grower (not observer), and I like "riding along" with the growing element to visualize what's happening. (Radians are similar: they measure angles in terms of the mover.)

Next question: what change do we apply the logarithm to?

We're really just interested in the ratio between start and finish: 9.9 trillion to 14.4 trillion in 10 years. This is the same growth *rate* as going from $9.90 to $14.40 in the same period.

We can sharpen our thought:

$$\text{natural logarithm of growth ratio} \rightarrow \text{cause of growth}$$

$$\ln\left(\frac{14.4}{9.9}\right) = .374$$

Ok, the cause was a rate of .374 or 37.4%. Are we done?

Not yet. Logarithms don't know about how long a change took (we didn't plug in 10 years, right?). They give us a rate as if all the change happened in a single time period.

The change could indeed be a single year of 37.4% continuous growth, or 2 years of 18.7% growth, or some other combination.

From the scenario, we know the change took 10 years, so the rate must have been:

$$\text{rate} = \frac{.374}{10} = .0374 = 3.74\%$$

From the viewpoint of instant, continuous growth, the US economy grew by 3.74% per year.

Are we done now? Not quite!

This continuous rate is from the grower's perspective, as if we're "riding along" with the economy as it changes. A banker probably cares about the human-friendly, year-over-year difference. We can figure this out by letting the continuous growth run for a year:

$$\text{exponent with rate \& time} \rightarrow \text{effect of growth}$$

$$e^{\text{rate}\cdot\text{time}} = \text{growth}$$

$$e^{.0374\cdot 1} = 1.0381$$

The year-over-year gain is 3.8%, slightly higher than the 3.74% instantaneous rate due to compounding. Here's another way to put it:

- From an instant-by-instant basis, a given part of the economy is growing by 3.74%, modeled by $e^{.0374 \cdot \text{years}}$
- On a year-by-year basis, with compounding effects worked out, the economy grows by 3.81%, modeled by 1.0381^{years}

In finance, we may want the year-over-year change which can be compared nicely with other trends. In science and engineering, we prefer modeling behavior on an instantaneous basis.

Scenario: Describing Natural Growth

I detest contrived examples like "Assume bacteria doubles every 24 hours, find its growth formula.". Do bacteria colonies replicate on clean human intervals, and do we wait around for an exact doubling?

A better scenario: "Hey, I found some bacteria, waited an hour, and the lump grew from 2.3 grams to 2.32 grams. I'm going to lunch now. Figure out how much we'll have when I'm back in 3 hours."

Let's model this. We'll need a logarithm to find the growth rate, and then an exponent to project that growth forward. Like before, let's keep everything in terms of the natural log to start.

The growth factor is:

$$\text{logarithm of change} \rightarrow \text{cause of growth}$$

$$\ln(\text{growth}) = \ln(2.32/2.3) = .0086 = .86\%$$

That's the rate for one hour, and the general model to project forward will be

$$\text{exponent with rate \& time} \rightarrow \text{effect of growth}$$

$$e^{.0086 \cdot \text{hours}} \rightarrow \text{effect of growth}$$

If we start with 2.32 and grow for 3 hours we'll have:

$$2.32 \cdot e^{.0086 \cdot 3} = 2.38$$

Just for fun, how long until the bacteria doubles? Imagine waiting for 1 to turn to 2:

$$1 \cdot e^{.0086 \cdot \text{hours}} = 2$$

We can mechanically take the natural log of both sides to "undo the exponent", but let's think intuitively.

If 2 is the final result, then ln(2) is the growth input that got us there (some rate × time). We know the rate was .0086, so the time to get to 2 would be:

$$\text{hours} = \frac{\ln(2)}{\text{rate}} = \frac{.693}{.0086} = 80.58$$

The colony will double after ~80 hours. (Glad you didn't stick around?)

What Does The Perspective Change Really Mean?

Figuring out whether you want the input (cause of growth) or output (result of growth) is pretty straightforward. But how do you visualize the grower's perspective?

Imagine we have little workers who are building the final growth pattern (see the [article on exponents]()):

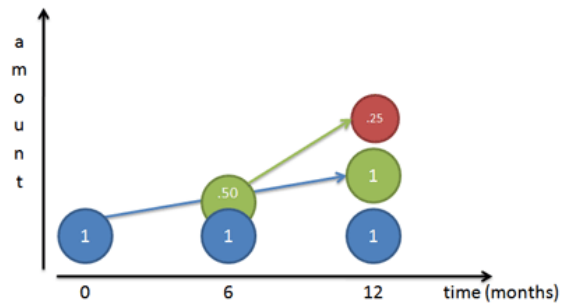

If our growth rate is 100%, we're telling our initial worker (Mr. Blue) to work steadily and create a 100% copy of himself by the end of the year. If we follow him day-by-day, we see he does finish a 100% copy of himself (Mr. Green) at the end of the year.

But… that worker he was building (Mr. Green) starts working as well. If Mr. Green first appears at the 6-month mark, he has a half-year to work (same annual rate as Mr. Blue) and he builds Mr. Red. Of course, Mr. Red ends up being half done, since Mr. Green only has 6 months.

What if Mr. Green showed up after 4 months? A month? A day? A second? If workers begin growing immediately, we get the instant-by-instant curve defined by e^x:

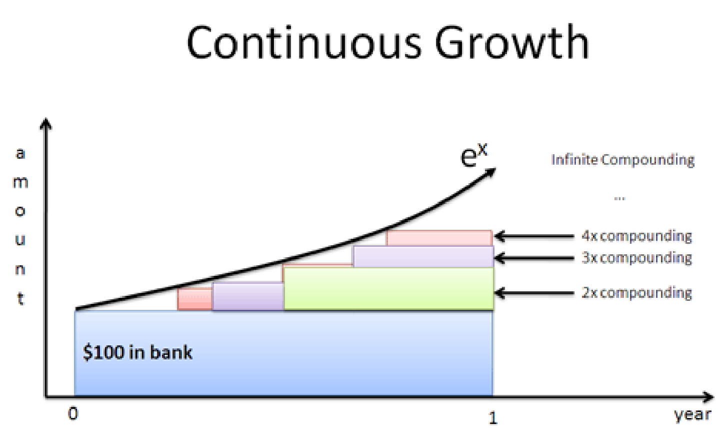

The natural log gives a growth rate in terms of an individual worker's perspective. We plug that rate into e^x to find the final result, with all compounding included.

Using Other Bases

Switching to another type of logarithm (base 10, base 2, etc.) means we're looking for some pattern in the overall growth, not what the individual worker is doing.

Each logarithm asks a question when seeing a change:

- Log base e: What was the instantaneous rate followed by each worker?
- Log base 2: How many doublings were required?
- Log base 10: How many 10x-ings were required?

Here's a scenario to analyze:

- Over 30 years, the transistor counts on typical chips went from 1000 to 1 billion

How would you analyze this?

- Microchips aren't a single entity that grow smoothly over time. They're separate editions, from competing companies, and indicate a general tech trend.
- Since we're not "riding along" with an expanding microchip, let's use a scale made for human convenience. Doubling is easier to think about than 10x-ing.

With these assumptions we get:

$$\text{logarithm of change} \rightarrow \text{cause of growth}$$

$$\log_2\left(\frac{1 \text{ billion}}{1000}\right) = \log_2(1 \text{ million}) \sim 20 \text{ doublings}$$

The "cause of growth" was 20 doublings, which we know occurred over 30 years. This averages 2/3 doublings per year, or 1.5 years per doubling — a nice rule of thumb.

From the grower's perspective, we'd compute $\frac{ln\left(\frac{1\ billion}{1000}\right)}{30} years = 46\%$ continuous growth (a bit harder to relate to in this scenario).

We can summarize our analysis in a table:

Scenario: 1000 to 1 billion transistors in 30 years

	Find Cause	Find Effect
From Grower's Viewpoint	$\ln\left(\frac{1 \text{ billion}}{1000}\right) = 13.81$ $\frac{13.81 \text{ "input cause"}}{30 \text{ years}} = .46 \text{ rate per year}$ *Find continuous growth rate*	$e^{0.46 \cdot \text{years}}$ *Model continuous growth*
From Observer's Viewpoint	$\log_2\left(\frac{1 \text{ billion}}{1000}\right) = 20 \text{ doublings}$ $\frac{20 \text{ doublings}}{30 \text{ years}} = \frac{2}{3} \text{ doublings per year}$ *Find doubling rate*	$2^{\frac{2}{3} \cdot \text{years}}$ *Model doubling rate*

Summary

Learning is about finding the hidden captions behind a concept. When is it used? What point view does it bring to the problem?

My current interpretation is that exponents ask about cause vs. effect and grower vs. observer. But we're never done; part of the fun is seeing how we can re-caption old concepts.

Happy math.

Appendix: The Change of Base Formula

Here's how to think about switching bases. Assuming a 100% continuous growth rate,

- ln(x) is the time to grow to x
- ln(2) is the time to grow to 2

Since we have the time to double, we can see how many would "fit" in the total time to grow to x:

$$\text{number of doublings from 1 to x} = \frac{\ln(x)}{\ln(2)} = \log_2(x)$$

For example, how many doublings occur from 1 to 64?

Well, $ln(64) = 4.158$. And $ln(2) = .693$. The number of doublings that fit is:

$$\frac{\ln(64)}{\ln(2)} = \frac{4.158}{.693} = 6$$

In the real world, calculators may lose precision, so use a direct log base 2 function if possible. And of course, we can have a fractional number: Getting from 1 to the square root of 2 is "half" a doubling, or $log_2(1.414) = 0.5$.

Changing to log base 10 means we're counting the number of 10x-ings that fit:

$$\text{number of 10x-ings from 1 to x} = \frac{\ln(x)}{\ln(10)} = \log_{10}(x)$$

Neat, right? Read Using Logarithms in the Real World for more examples.