6

On Number Numbness

May, 1982

THE renowned cosmogonist Professor Bignumska, lecturing on the future of the universe, had just stated that in about a billion years, according to her calculations, the earth would fall into the sun in a fiery death. In the back of the auditorium a tremulous voice piped up: "Excuse me, Professor, but h-h-how long did you say it would be?" Professor Bignumska calmly replied, "About a billion years." A sigh of relief was heard. "Whew! For a minute there, I thought you'd said a *million* years."

John F. Kennedy enjoyed relating the following anecdote about a famous French soldier, Marshal Lyautey. One day the marshal asked his gardener to plant a row of trees of a certain rare variety in his garden the next morning. The gardener said he would gladly do so, but he cautioned the marshal that trees of this size take a century to grow to full size. "In that case," replied Lyautey, "plant them this afternoon."

In both of these stories, a time in the distant future is related to a time closer at hand in a startling manner. In the second story, we think to ourselves: Over a century, what possible difference could a day make? And yet we are charmed by the marshal's sense of urgency. Every day counts, he seems to be saying, and particularly so when there are thousands and thousands of them. I have always loved this story, but the other one, when I first heard it a few thousand days ago, struck me as uproarious. The idea that one could take such large numbers so personally, that one could sense doomsday so much more clearly if it were a mere *million* years away rather than a far-off *billion* years—hilarious! Who could possibly have such a gut-level reaction to the difference between two huge numbers?

Recently, though, there have been some even funnier big-number "jokes" in newspaper headlines—jokes such as "Defense spending over the next four years will be $1 trillion" or "Defense Department overrun over the next four years estimated at $750 billion". The only thing that worries me about these jokes is that their humor probably goes unnoticed by the average citizen. It would be a pity to allow such mirth-provoking notions to be appreciated only by a select few, so I decided it would be a good idea to devote some space to the requisite background knowledge, which also happens to be one of my favorite topics: the lore of very large (and very small) numbers.

I have always suspected that relatively few people really know the difference between a million and a billion. To be sure, people generally know it well enough to sense the humor in the joke about when the earth will fall into the sun, but what the difference is *precisely*—well, that is something else. I once heard a radio news announcer say, "The drought has cost California agriculture somewhere between nine hundred thousand and a billion dollars." Come again? This kind of thing worries me. In a society where big numbers are commonplace, we cannot afford to have such appalling number ignorance as we do. Or do we actually suffer from *number numbness?* Are we growing ever number to ever-growing numbers?

What do people think when they read ominous headlines like the ones above? What do they think when they read about nuclear weapons with 20-kiloton yields? Or 60-megaton yields? Does the number really register —or is it just another cause for a yawn? "Ho hum, I always knew the Russians could kill us all 20 times over. So now it's 200 times, eh? Well, we can be thankful it's not 2,000, can't we?"

What do people think about the fact that in some heavily populated areas of the U.S., it is typical for the price of a house to be a quarter of a million dollars? What do people think when they hear radio commercials for savings institutions telling them that if they invest now, they could have a million dollars on retirement? Can *everyone* be a millionaire? Do we now *expect* houses to take a fourth of a millionaire's fortune? What ever has become of the once-glittery connotations of the word "millionaire"?

* * *

I once taught a small beginning physics class on the thirteenth floor of Hunter College in New York City. From the window we had a magnificent view of the skyscrapers of midtown Manhattan. In one of the opening sessions, I wanted to teach my students about estimates and significant figures, so I asked them to estimate the height of the Empire State Building. In a class of ten students, not one came within a factor of two of the correct answer (1,472 feet with the television antenna, 1,250 without). Most of the estimates were between 300 and 500 feet. One person thought 50 feet was right—a truly amazing underestimate; another thought it was a mile. It turned out that this person had actually calculated the answer, guessing 50 feet per story and 100 stories or so, thus getting about 5,000 feet. Where one person thought each *story* was 50 feet high, another thought the whole 102-story *building* was that high. This startling episode had a deep effect on me.

It is fashionable for people to decry the appalling illiteracy of this generation, particularly its supposed inability to write grammatical English. But what of the appalling *innumeracy* of most people, old and young, when it comes to making sense of the numbers that, in point of fact, and whether they like it or not, run their lives? As Senator Everett Dirksen once said, "A billion here, a billion there—soon you're talking real money."

The world is gigantic, no question about it. There are a lot of people, a lot of needs, and it all adds up to a certain degree of incomprehensibility. But that is no excuse for not being able to understand—or even relate to —numbers whose purpose is to summarize in a few symbols some salient aspects of those huge realities. Most likely the readers of this article are not the ones I am worried about. It is nonetheless certain that every reader of this article knows many people who are ill at ease with large numbers of the sort that appear in our government's budget, in the gross national product, corporation budgets, and so on. To people whose minds go blank when they hear something ending in "illion", all big numbers are the same, so that exponential explosions make no difference. Such an inability to relate to large numbers is clearly bad for society. It leads people to ignore big issues on the grounds that they are incomprehensible. The way I see it, therefore, anything that can be done to correct the rampant innumeracy of our society is well worth doing. As I said above, I do not expect this article to reveal profound new insights to its readers (although I hope it will intrigue them); rather, I hope it will give them the materials and the impetus to convey a vivid sense of numbers to their friends and students.

* * *

As an aid to numerical horse sense, I thought I would indulge in a small orgy of questions and answers. Ready? Let's go! How many letters are there in a bookstore? Don't calculate—just guess. Did you say about a billion? That has nine zeros (1,000,000,000). If you did, that is a pretty sensible estimate. If you didn't, were you too high or too low? In retrospect, does your estimate seem far-fetched? What intuitive cues suggest that a billion is appropriate, rather than, say, a million or a trillion? Well, let's calculate it. Say there are 10,000 books in a typical bookstore. (Where did I get this? I just estimated it off the top of my head, but on calculation, it seems reasonable to me, perhaps a bit on the low side.) Now each book has a couple of hundred pages filled with text. How many words per page—a hundred? A thousand? Somewhere in between, undoubtedly. Let's just say 500. And how many letters per word? Oh, about five, on the average. So we have $10{,}000\times200\times500\times5$, which comes to five billion. Oh, well—who cares about a factor of five when you're up this high? I'd say that if you were within a factor of ten of this (say, between 500 million and 50 billion), you were doing pretty well. Now, could we have sensed this *in advance*—by which I mean, *without calculation?*

We were faced with a choice. Which of the following twelve possibilities is the most likely:

(a) 10;
(b) 100;
(c) 1,000;
(d) 10,000;
(e) 100,000;
(f) 1,000,000;
(g) 10,000,000;
(h) 100,000,000;
(i) 1,000,000,000;
(j) 10,000,000,000;
(k) 100,000,000,000;
(l) 1,000,000,000,000?

In the United States, this last number, with its twelve zeros, is called a *trillion;* in most other countries it is called a *billion.* People in those countries reserve "trillion" for the truly enormous number 1,000,000,000,000,000,000—to us a "quintillion"—though hardly anyone knows that term.

What most people truly don't appreciate is that making such a guess is very much the same as looking at the chairs in a room and guessing quickly if there are two or seven or fifteen. It is just that here, what we are guessing at is the number of zeros in a numeral, that is, the logarithm (to the base 10) of the number. *If we can develop a sense for the number of chairs in a room, why not as good a sense for the number of zeros in a numeral?* That is the basic premise of this article.

Of course there is a difference between these two types of numerical horse sense. It is one thing to look at a numeral such as "1000000000000" and to have an intuitive feeling, without counting, that it has somewhere around twelve zeros—certainly more than ten and fewer than fifteen. It is quite another thing to look at an aerial photograph of a logjam (see Figure 6-1) and to be able to sense, visually or intuitively or somewhere in between, that there must be between three and five zeros in the decimal representation of the number of logs in the jam—in other words, that 10,000 is the closest power of 10, that 1,000 would definitely be too low, and that 100,000 would be too high. Such an ability is simply a form of number perception one level of abstraction higher than the usual kind of number perception. But one level of abstraction should not be too hard to handle.

The trick, of course, is practice. You have to get used to the idea that ten is a very big number of zeros for a numeral to have, that five is pretty big, and that three is almost graspable. Probably what is most important is that you should have a prototype example for each number of zeros. For instance: *Three* zeros would take care of the number of students in your high school: 1,000, give or take a factor of three. (In numbers having just a few zeros we are always willing to forgive a factor of three or so in either direction, as long as we are merely estimating and not going for exactness.) *Four* zeros is the number of books in a non-huge bookstore. *Five* zeros is

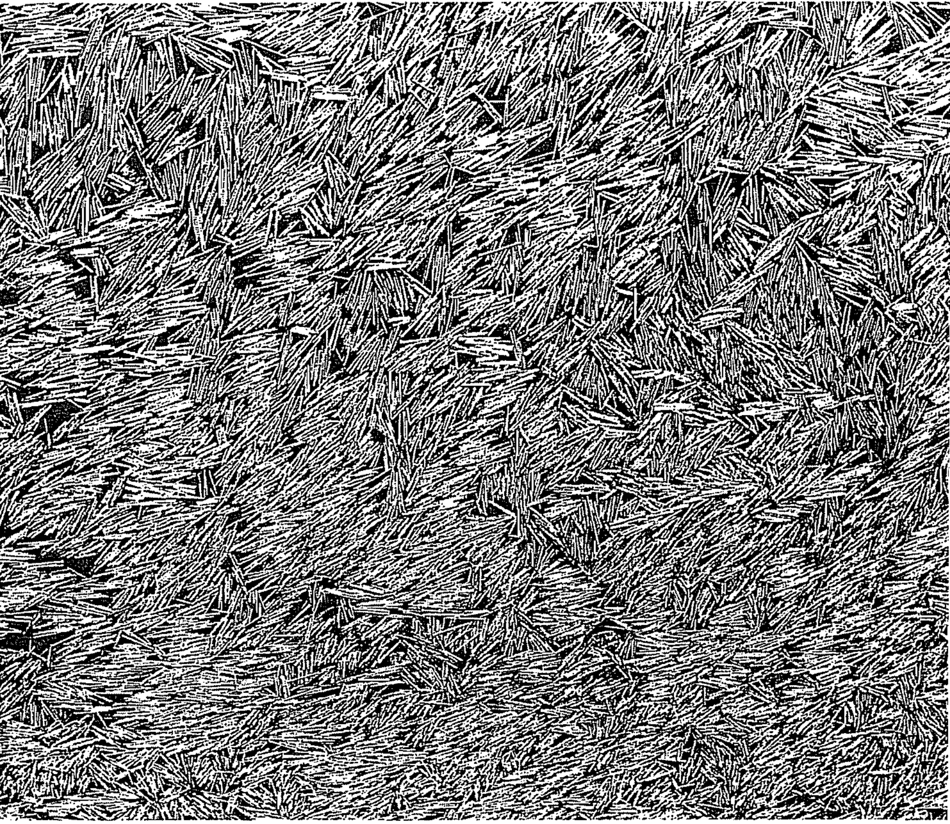

FIGURE 6–1. *Aerial view of a logjam in Oregon. How many logs? [Photo by Ray Atkeson.]*

the size of a typical county seat: 100,000 souls or so. *Six* zeros—that is, a million—is getting to be a large city: Minneapolis, San Diego, Brasília, Marseilles, Dar és Salaam. *Seven* zeros is getting huge: Shanghai, Mexico City, Seoul, Paris, New York. Just how many cities do you think there are in the world with a population of a million or more? Of them, how many do you think you have never heard of? What if you lowered the threshold to 100,000? How many towns are there in the United States with a population of 1,000 or less? Here is where practice helps.

I said that you should have one prototype example for each number of digits. Actually, that is silly. You should have a few. In order to have a concrete sense of "nine-zero-ness", you need to see it instantiated in several different media, preferably as diverse as populations, budgets, small objects (ants, coins, letters, etc.), and maybe a couple of miscellaneous places, such as astronomical distances or computer statistics.

Consider the famous claim made by the McDonald's hamburger chain: "Over 25 billion served" (or whatever they say these days). Is this figure credible? Well, if it were ten times bigger—that is, 250 billion—we could divide by the U.S. population more easily. (This is apparent if you happen to know that the U.S. population is about 230 million. For the purposes of this discussion, let us call the U.S. population 250 million, or 2.5×10^8—a common number that everyone should know.) Let us imagine, then, that the claim were "Over 250 billion served". Then we would compute that 1,000 burgers had been cooked for every person in the U.S. But since we deliberately inflated it by a factor of 10, let us now undo that—let us divide our answer by ten, to get 100. Is it plausible that McDonald's has prepared 100 burgers for every person in the U.S.? Sounds reasonable to me; after all, they have been around for many years, and some families go there many times a year. Therefore the claim *is* plausible, and the fact that it is *plausible* makes it *probable* that it is quite accurate. Presumably, McDonald's wouldn't go to the trouble of updating their signs every so often if they were not trying to be accurate. I must say that if their earnest effort helps to reduce innumeracy, I approve highly of it.

Where do all those burgers come from? A staggering figure is the number of cattle slaughtered every day in the U.S. It comes to about 90,000. When I first heard this, it sounded amazingly high, but think about it. Maybe half a pound of meat per person per day. Once again, the U.S. population—250 million—comes in handy. With half a pound of meat per person per day, that comes to 100 million pounds of meat per day—or something like that, anyway. We're certainly not going to worry about factors of two. How many tons is that? Divide by 2,000 to get 50,000 tons. But an individual animal does not yield a ton of meat. Maybe 1,000 pounds or so—half a ton. For each ton of meat, that would mean two animals were killed. So we would get about 100,000 animals biting the dust every day to satisfy our collective appetite. Of course, we do not eat only beef, so the true figure should be a bit lower. And that brings us back down to about the right figure.

* * *

How many trees are cut down each week to produce the Sunday edition of the *New York Times*? Say a couple of million copies are printed, each one weighing four pounds. That comes to about eight million pounds of paper —4,000 tons. If a tree yielded a ton of paper, that would be 4,000 trees. I don't know much about logging, but we cannot be too far off in assuming a ton per tree. At worst it would be 200 pounds of paper per tree, and that would mean 40,000 small trees. The logjam photograph shows somewhere between 7,500 and 15,000 logs, as nearly as I can estimate. So, if we do assume 200 pounds of paper per tree, the logs in the photograph represent considerably less than half of one Sunday *Times*' worth of trees! We could go on to estimate the number of trees cut down every month to provide for all the magazines, books, and newspapers published in this country, but I'll leave that to you.

How many cigarettes are smoked in the U.S. every year? (How many zeros?) This is a classic "twelver"—on the order of a trillion. It is easy to calculate. Say that half of the people in the country are cigarette smokers: 100 million of them. (I know this is something of an overestimate; we'll compensate by reducing something else somewhere along the way.) Each smoker smokes—what? A pack per day? All right. That makes 20 cigarettes times 100 million: two billion cigarettes per day. There are 365 days per year, but let's say 250, since I promised to reduce something somewhere; 250 times two billion gives about 500 billion—half a trillion. This is just about on the nose, as it turns out; the last I looked (a few years ago), it was some 545 billion. I remember how awed I was when I first encountered this figure; it was the first time I had met up with a *concrete* number about the size of a trillion.

By the way, "20 (cigarettes) times 100 million" is not a hard calculation, yet I bet it would stump many Americans, if they had to do it in their head. My way of doing it is to shift a factor of 10 from one number to the other. Here, I *reduce* 20 to 2, while *increasing* 100 to 1,000. It makes the problem into "2 times 1,000 million", and then I just remember that 1,000 million is one billion. I realize that this sounds absolutely trivial to anyone who is comfortable with figures, but it sounds truly frightening and abstruse to people who are not so comfortable with them—and that means *most* people.

It is numbers like 545 billion that we are dealing with when we talk about a Defense Department overrun of $750 billion for the next four years. A really fancy single-user computer (the kind I wouldn't mind having) costs approximately $75,000. With $750 billion to throw around, we could give one to every person in New York City, which is to say, we could buy about ten million of them. Or, we could give $1 million to every person in San Francisco, and still have enough left over to buy a bicycle for everyone in China! There's no telling what good uses we could put $750 billion to. But instead, it will go into bullets and tanks and fighters and war games and missile systems and jet fuel and marching bands and so on. An interesting way to spend $750 billion, but I can think of better ways.

* * *

Let us think of some other kinds of big numbers. Did you know that your retina has about 100 million cells in it, each of which responds to some particular kind of stimulus? And they feed their signals back into your brain, which is now thought to consist of somewhere around 100 billion neurons, or nerve cells. The number of glia—smaller supporting cells in the brain—is about ten times as large. That means you have about one trillion glia in your little noggin. That may sound big, but in your body altogether there are estimated to be about 60 or 70 trillion cells. Each one of them contains millions of components working together. Take the protein hemoglobin, for instance, which transports oxygen in the bloodstream. We each have about six billion trillion (that is, six thousand million million million) copies of the hemoglobin molecule inside us, with something like 400 trillion of them (400 million million) being destroyed every second, and another 400 trillion being made! (By the way, I got these figures from Richard Dawkins' book *The Selfish Gene.* They astounded me when I read them there, and so I tried to calculate them on my own. My estimates came out pretty close to his figures, and then, for good measure, I asked a friend in biology to calculate them, and she seemed to get about the same answers independently, so I guess they are pretty reliable.)

The number of hemoglobin molecules in the body is about 6×10^{21}. It is a curious fact that over the past year or two, nearly everyone has become familiar, implicitly or explicitly, with a number nearly as big—namely, the number of different possible configurations of Rubik's Cube. This number —let us call it *Rubik's constant*—is about 4.3×10^{19}. For a very vivid image of how big this is, imagine that you have many cubes, an inch on each side, one in every possible configuration. Now you start spreading them out over the surface of the United States. How thickly covered would the U.S. be in cubes? Moreover, if you are working in Rubik's "supergroup", where the orientations of face centers matter, then Rubik's "superconstant" is 2,048 times bigger, or about 9×10^{22}!

The Ideal Toy Corporation—American marketer of the Cube—was far less daring than McDonald's. On their package, they softened the blow, saying merely "Over three billion combinations possible"—a pathetic and euphemistic underestimate if ever I heard one. This is the first time I have ever heard Muzak based on a pop number rather than a pop melody. Try these out, for comparison's sake:

(1) "Entering San Francisco—population greater than 1."
(2) "McDonald's—over 2 served."
(3) "Together, the superpowers have 3 pounds of TNT for every human being on earth."

Number 1 is off by a factor of about a million, or six *orders of magnitude* (factors of ten). Number 2 is off by a factor of ten billion or so (ten orders of magnitude), while number 3 (which I saw in a recent letter to the editor of the *Bulletin of the Atomic Scientists*) is too small by a factor of about a thousand (three orders of magnitude).

The hemoglobin number and Rubik's superconstant are *really* big. How about some smaller big ones, to come back to earth for a moment? All right —how many people would you say are falling to earth by parachute at this moment (a perfectly typical moment, presumably)? How many English words do you know? How many murders are there in Los Angeles County every year? In Japan? These last two give quite a shock when put side by side: Los Angeles County, about 2,000; Japan, about 900.

Speaking of yearly deaths, here is one we are all used to sweeping under the rug, it seems: 50,000 dead per year (in this country alone) in car accidents. If you count the entire world, it's probably two or three times that many. Can you imagine how we would react if someone said to us today: "Hey, everybody! I've come up with a really nifty invention. Unfortunately, it has a minor defect—every twelve years or so it will wipe out about as many Americans as the population of San Francisco. But wait a minute! Don't go away! The rest of you will love it, I promise!" Now, these statistics are accurate for cars. And yet we seldom hear people chanting, "No cars is good cars!" How many bumper strips have you seen that say, "No more cars!"? Somehow, collectively, we are willing to absorb the loss of 50,000 lives per year without any serious worry. And imagine that half of this—25,000 needless deaths—is due to drunks behind the wheel. Why aren't you just fuming?

*   *   *

I said I would be a little lighter. All right. Light consists of photons. How many photons per second does a 100-watt bulb put out? About 10^{20}—another biggie. Is it bigger or smaller than the number of grains of sand on a beach? What beach? Say a stretch of beach a mile long, 100 feet wide and six feet deep. What would you estimate? Now calculate it. How about trying the number of drops in the Atlantic Ocean? Then try the number of fish in the ocean. Which are there more of: fish in the sea, or ants on the surface of the earth? Atoms in a blade of grass, or blades of grass on the earth? Blades of grass, or insects? Leaves on a typical oak tree, or hairs on a human head? How many raindrops fall on your town in one second during a terrific downpour?

How many copies of the Mona Lisa have ever been printed? Let's try this one together. Probably it is printed in magazines in the United States a few dozen times per year. Say each of the magazines prints 100,000 copies. That makes a few million copies per year in American magazines, but then there are books and other publications. Maybe we should double or triple our figure for the U.S. To take into account other countries, we can multiply it again by three or four. Now we have hit about 100 million copies per year. Let us assume this held true for each year of this century. That would make nearly ten billion copies of the Mona Lisa! Quite a meme, eh? Probably we have made some mistakes along the way, but give or take a factor of ten, that is very likely about what the number is.

"Give or take a factor of *ten*"!? A moment ago I was saying that a factor of *three* was forgivable, but now, here I am forgiving myself *two* factors of three—that is, an entire order of magnitude. Well, the reason is simple: We are now dealing with larger numbers (10^{10} instead of 10^5), and so it is permissible. This brings up a good rule of thumb. Say an error of a factor of three is permissible for each estimated factor of 100,000. That means we are allowed to be off by a factor of ten—*one* order of magnitude—when we get up to sizes around ten billion, or by a factor of 100 or so (*two* orders of magnitude) when we get up to the square of that, which is 10^{20}, about 2.5

times the size of Rubik's constant. This means it would have been forgivable if Ideal had said, "Over a *billion* billion combinations", since then they would have been off by a factor of only 40—about 1.5 orders of magnitude —which is within our limits when we're dealing with numbers that large.

Why should we be content with an estimate that is only one percent of the actual number, or with an estimate that is 100 times too big? Well, if you consider the base-10 logarithm of the number—the number of zeros—then if we say 18 when the real answer is 20, we are off by only ten percent! Now what entitles us to cavalierly dismiss the magnitude itself and to switch our focus to its logarithm (its order of magnitude)? Well, when numbers get this big, we have no choice. Our perceptual reality begins to shift. We simply *cannot* visualize the actual quantity. The numeral—the string of digits—takes over: our perceptual reality becomes one of numbers of zeros. When does this shift take place? It begins when we can no longer see, in our mind's eye, a collection of the right order of magnitude. For me, this "perceptual logjam" begins at about 10^4—the size of the actual logjam I remember in the photograph. It is important to understand this transition. It is one of the key ideas of this article.

There are other ways to grasp 10^4, such as the number of soup cans that would fill a 50-foot shelf in a supermarket. Numbers much bigger than that, I simply cannot visualize. The number of tiles lining the Lincoln Tunnel between Manhattan and New Jersey is so enormous that I cannot easily picture it. (It is on the order of a million, as you can calculate for yourself, even if you've never seen it!) In any case, somewhere around 10^4 or 10^5, my ability to visualize begins to fade and to be replaced with that second-order reality of the number of digits (or, to some extent, with number names such as "million", "billion", and "trillion"). Why it happens at this size and not, say, at 10 million or at 1,000 must have to do with evolution and the role that the perception of vast arrays plays in survival. It is a fascinating philosophical question, but one I cannot hope to answer here.

In any case, a pretty good rule of thumb is this: Your estimate should be within ten percent of the correct answer—but this need apply only *at the level of your perceptual reality.* Therefore you are excused if you guessed that Rubik's cube has 10^{18} positions, since 18 is pretty close to 19.5, which is about what the number of digits is. (Remember that—roughly speaking—Rubik's constant is 4.3×10^{19}, or 43,000,000,000,000,000,000. The leading factor of 4.3 counts for a bit more than half a digit, since each factor of 10 contributes a *full* digit, whereas a factor of 3.16, the square root of 10, contributes *half* a digit.)

If, perchance, you were to start dealing with numbers having millions or billions of digits, the numerals themselves (the colossal strings of digits) would cease to be visualizable, and your perceptual reality would be forced to take another leap upward in abstraction—to the number that counts the digits in the number that counts the digits in the number that counts the objects concerned. Needless to say, such third-order perceptual reality is highly abstract. Moreover, it occurs very seldom, even in mathematics. Still, you can imagine going far beyond it. Fourth- and fifth-order perceptual realities would quickly yield, in our purely abstract imagination, to tenth-, hundredth-, and millionth-order perceptual realities.

By this time, of course, we would have lost track of the *exact* number of levels we had shifted, and we would be content with a mere *estimate* of that number (accurate to within ten percent, of course). "Oh, I'd say about two million levels of perceptual shift were involved here, give or take a couple of hundred thousand" would be a typical comment for someone dealing with such unimaginably unimaginable quantities. You can see where this is leading: to multiple levels of abstraction in talking about multiple levels of abstraction. If we were to continue our discussion just one zillisecond longer, we would find ourselves smack-dab in the middle of the theory of recursive functions and algorithmic complexity, and that would be too abstract. So let's drop the topic right here.

* * *

Related to this idea of huge numbers of digits, but more tangible, is the computation of the famous constant π. How many digits have so far been calculated by machine? The answer (as far as I know) is one million. It was done in France a few years ago, and the million digits fill an entire book. Of these million, how many have been committed to human memory? The answer strains credulity: 20,000, according to the latest *Guinness Book of World Records.* I myself once learned 380 digits of π, when I was a crazy high-school kid. My never-attained ambition was to reach the spot, 762 digits out in the decimal expansion, where it goes "999999", so that I could recite it out loud, come to those six '9's, and then impishly say, "and so on!" Later, I met several other people who had outdone me (although none of them had reached that string of '9's). All of us had forgotten most of the digits we once knew, but at least we all remembered the first 100 solidly, and so occasionally we would recite them in unison—a rather esoteric pleasure.

What would you think if someone claimed that the entire book of a million digits of π had been memorized by someone? I would dismiss the claim out of hand. A student of mine once told me very earnestly that Jerry Lucas, the memory and basketball whiz, knew the entire Manhattan telephone directory by heart. Here we have a good example of how innumeracy can breed gullibility. Can you imagine what memorizing the Manhattan telephone directory would involve? To me, it seems about two orders of magnitude beyond credibility. To memorize one page seems fabulously difficult. To memorize ten pages seems at about the limit of credibility. Incidentally, memorizing the entire Bible (which I have occasionally heard claimed) seems to me about equivalent to memorizing ten pages of the phone book, because of the high redundancy of written language and the regularity of events in the world. But to have memorized 1,500 dense pages of telephone numbers, addresses, and names is literally beyond belief. I'll eat my hat—in fact, all of my 10,000 hats—if I'm wrong.

* * *

There are some phenomena for which there are two (or more) scales with which we are equally comfortable, depending on the circumstances. Take pitch in music. If you look at a piano keyboard, you will see a linear scale along which pitch can be measured. The natural thing to say is: "This A is nine semitones higher than that C, and the C is seven semitones higher than that F, so the A is 16 semitones higher than the F." It is an additive, or linear, scale. By this I mean that if you assigned successive whole numbers to successive notes, then the distance from any note to any other would be given by the difference between their numbers. Only addition and subtraction are involved.

By contrast, if you are going to think of things acoustically rather than auditorily, physically rather than perceptually, each pitch is better described in terms of its *frequency* than in terms of its position on a keyboard. The low A at the bottom of the keyboard vibrates about 27 times per second, whereas the C three semitones above it vibrates about 32 times per second. So you might be inclined to guess that in order to jump up three semitones one should always add five cycles per second. Not so. You should always *multiply* by about 32/27 instead. If you jump up twelve semitones, that means four repeated up-jumps of three semitones.

Thus, when you have gone up one octave (twelve semitones), your pitch has been multiplied by 32/27 four times in a row, which is 2. Actually, the fourth power of 32/27 is not quite 2, and since an octave represents a ratio of *exactly* 2, 32/27 must be a slight underestimate. But that is beside the point. The point is that the natural operations for comparing frequencies are multiplication and division, whereas the natural operations for note numbers on a keyboard are addition and subtraction. What this means is that the note numbers are logarithms of the frequencies. Here is a case where we think naturally in logarithms!

Here is a different way of putting things. Two adjacent notes near the top of a piano keyboard differ in frequency by about 400 cycles per second, whereas adjacent notes near the bottom differ by only about two cycles per second. Wouldn't that seem to imply that the intervals are wildly different? Yet to the human ear, the high and the low interval sound exactly the same!

Logarithmic thinking happens when you perceive only a linear increase even if the thing itself doubles in size. For instance, have you ever marveled at the fact that dialing a mere seven digits can connect any telephone to any other in the New York metropolitan area, where some 10 million people live? Suppose New York were to double in population. Would you then have to add seven more digits to each phone number, making fourteen-digit numbers, in order to reach those twenty million people? Of course not.

Adding seven more digits would *multiply* the number of possibilities by ten million. In fact, adding merely three digits (the area code in front) enables you to reach any phone number in North America. This is simply because each new digit creates a tenfold increase in the number of phones reachable. Three more digits will always multiply your network by a factor of 1,000: three orders of magnitude. Thus the length of a phone number—the quantity directly perceived by you when you are annoyed at how long it takes to dial a long-distance number—is a logarithmic measure of the size of the network you are embedded in. That is why it is preposterous to see huge long numbers of 25 or 30 digits used as codes for people or products when, without any doubt, a few digits would suffice.

I once was sent a bill asking that I transfer a fee to account No. 60802-620-1-1-721000-421-01062 in a bank in Yugoslavia. For a while this held my personal record for absurdity of numbers encountered in business transactions. Recently, however, I was sent my car registration form, at the bottom of which I found this enlightening constant: 01010136121820030107001426311724151200360360003000 2. For good measure it was followed, a few blank spaces later, by '19283'.

One place where we think logarithmically is number names. We in America have a new name every three zeros (up to a certain point): from *thousand* to *million* to *billion* to *trillion*. Each jump is "the same size", in a sense. That is, a billion is exactly as much bigger than a million as a million is bigger than a thousand. Or a trillion is to a billion exactly as a billion is to a million. On the other hand, does this continue forever? For instance, does it seem reasonable to say that 10^{103} is to 10^{100} exactly as a million is to a thousand? I would be inclined to say "No, those big numbers are almost the same size, whereas a thousand and a million are very different." It is a little tricky because of the shifts in perceptual reality.

In any case, we seem to run out of number names at about a trillion. To be sure, there are some official names for bigger numbers, but they are about as familiar as the names of extinct dinosaurs: "quadrillion", "octillion", "vigintillion", "brontosillion", "triceratillion", and so on. We are simply not familiar with them, since they died off a dinosillion years ago. Even "billion" presents cross-cultural problems, as I mentioned above. Can you imagine what it would be like if in Britain, "hundred" meant 1,000? The fact is that when numbers get too large, people's imaginations balk. It is too bad, though, that a trillion is the largest number with a common name. What is going to happen when the defense budget gets even more bloated? Will we just get number? Of course, like the dinosaurs, we may never be granted the luxury of facing that problem.

*   *   *

The speed of automatic computation is something whose progress is best charted logarithmically. Over the past several decades, the number of primitive operations (such as addition or multiplication) that a computer can carry out per second has multiplied tenfold about every seven years. Nowadays, it is some 100 million operations per second or, on the fanciest machines, a little more. Around 1975, it was about 10 million operations per second. In the later 1960's, one million operations per second was extremely fast. In the early 1960's, it was 100,000 operations per second. 10,000 was high in the mid-1950's, 1,000 in the late 1940's—and in the early 1940's, 100.

In fact, in the early 1940's, Nicholas Fattu was the leader of a team at the University of Minnesota that was working for the Army Air Force on some statistical calculations involving large matrices (about 60×60). He brought about ten people together in a room, each of whom was given a Monroematic desk calculator. These people worked full-time for ten months in a coordinated way, carrying out the computations and cross-checking each other's results as they went along. About twenty years later, out of curiosity, Professor Fattu redid the calculations on an IBM 704 in twenty minutes. He found that the original team had made two inconsequential errors. Nowadays, of course, the whole thing could be done on a big "mainframe" computer in a second or two.

Still, modern computers can easily be pushed to their limits. The notorious computer proof of the four-color theorem, done at the University of Illinois a few years ago, took 1,200 hours of computer time. When you convert that into days, it sounds more impressive: 50 full 24-hour days. If the computer was carrying out twenty million operations per second, that would come to 10^{14}, or 100 trillion, primitive operations—a couple of hundred for every cigarette smoked that year in the U.S. Whew!

A computer doing a billion operations per second would really be moving along. Imagine breaking up one second into as many tiny fragments as there are seconds in 30 years. That is how tiny a nanosecond—a billionth of a second—is. To a computer, a second is a lifetime! Of course, the computer is dawdling compared with the events inside the atoms that compose it. Take one atom. A typical electron circling a typical nucleus makes about 10^{15} orbits per second, which is to say, a million orbits per nanosecond. From an electron's-eye point of view, a computer is as slow as molasses in January.

Actually, an electron has two eyes with which to view the situation. It has both an *orbital* cycle time and a *rotational* cycle time, since it is spinning on its own axis. Now, strictly speaking, "spin" is just a metaphor at the quantum level, so you should take the following with a big grain of salt. Nevertheless, if you imagine an electron to be a classically (non-quantum-mechanically) spinning sphere, you can calculate its rotation time from its known spin angular momentum (which is about Planck's constant, or 10^{-34} joule-second) and its radius (which we can equate with its Compton wavelength, which is about 10^{-10} centimeter). The spin time turns out to be about 10^{-20} second. In other words, every time the superfast computer adds two numbers, every electron inside it has pirouetted on its own axis about

100 billion times. (If we took the so-called "classical radius" of the electron instead, we would have the electron spinning at about 10^{24} times per second —enough to make one dizzy! Since this figure violates both relativity and quantum mechanics, however, let us be content with the first figure.)

At the other end of the scale, there is the slow, stately twirling of our galaxy, which makes a leisurely complete turn every 200 million years or so. And within the solar system, the planet Pluto takes about 250 years to complete an orbit of the sun. Speaking of the sun, it is about a million miles across and has a mass on the order of 10^{30} kilograms. The earth is a featherweight in comparison, a mere 10^{24} kilograms. And we should not forget that there are some stars—red giants—of such great diameter that they would engulf the orbit of Jupiter. Of course, such stars are very tenuous, something like cotton candy on a cosmic scale. By contrast, some stars—neutron stars—are so tightly packed that if you could remove from any of them a cube a millimeter on an edge, its mass would be about half a million tons, equal to the mass of the heaviest oil tanker ever built, fully loaded!

*   *   *

These large and small numbers are so far beyond our ordinary comprehension that it is virtually impossible to keep on being more amazed. The numbers are genuinely beyond understanding—unless one has developed a vivid feeling for various exponents. And even with such an intuition, it is hard to give the universe its awesome due for being so extraordinarily huge and at the same time so extraordinarily fine-grained. Number numbness sets in early these days. Most people seem entirely unfazed by words such as "billion" and "trillion"; they simply become synonyms for the meaningless "zillion".

This hit me particularly hard a few minutes after I had finished a draft of this column. I was reading the paper, and I came across an article on the subject of nerve gas. It stated that President Reagan expected the expenditures for nerve gas to come to about $800 million in 1983, and $1.4 billion in 1984. I was upset, but I caught myself being thankful that it was not $10 billion or $100 billion. Then, all at once, I really felt ashamed of myself. That guy has some nerve gas! How could I have been *relieved* by the figure of a "mere" $1.4 billion? How could my thoughts have become so dissociated from the underlying reality? One billion for nerve gas is not merely lamentable; it is odious. We cannot afford to become number-number than we are. We need to be willing to be jerked out of our apathy, because this kind of "joke" is in very poor taste.

Survival of our species is the name of the game. I don't really care if the number of mosquitoes in Africa is greater or less than the number of pennies in the gross national product. I don't care if there are more glaciers in the Dead Sea or scorpions in Antarctica. I don't care how tall a stack of one billion dollar bills would be (an image that President Reagan evoked in a speech decrying the size of the national debt created by his predecessors). I don't care a hoot about pointless, silly images of colossal magnitudes. What I *do* care about is what a billion dollars *represents* in terms of buying power: lunches for all the schoolkids in New York for a year, a hundred libraries, fifty jumbo jets, a few years' budget for a large university, one battleship, and so on. Still, if you love numbers (as I do), you can't help but blur the line between number play and serious thinking, because a silly image converts into a more serious image quite fluidly. But frivolous number virtuosity, enjoyable though it is, is far from the point of this article.

What I hope people will get out of this article is not a few amusing tidbits for the next cocktail party, but an increased passion about the importance of grasping large numbers. I want people to understand the very real consequences of those very surreal numbers bandied about in the newspaper headlines as interchangeably as movie stars' names in the scandal sheets. *That*'s the only reason for bringing up all the more humorous examples. At bottom, we are dealing with perceptual questions, but ones with life-and-death consequences!

*   *   *

Combatting number numbness is basically not so hard. It simply involves getting used to a second set of meanings for *small* numbers—namely, the meanings of numbers between say, five and twenty, when used as exponents. It would seem revolutionary for newspapers to adopt the convention of expressing large numbers as powers of ten, yet to know that a number has twelve zeros is *more* concrete than to know that it is called a "trillion".

I wonder what percentage of our population, if shown the numerals "314,159,265,358,979" and "271,828,182,845", would recognize that the former magnitude is about 1,000 times greater than the latter. I am afraid that the vast majority would not see it and would not even be able to read these numbers out loud. If that is the case, it is something to be worried about.

One book that attempts valiantly and poetically to combat such numbness, a book filled with humility before some of the astounding magnitudes that we have been discussing, is called *Cosmic View: The Universe in Forty Jumps,* by a Dutch schoolteacher, the late Kees Boeke. In his book, Boeke takes us on an imaginary voyage in pictures, in which each step is an exponential one, involving a factor of ten in linear size. From our own size, there are 26 upward steps and 13 downward steps. It is probably not coincidental that the book was written by someone from Holland, since the Dutch have long been internationally minded, living as they do in a small and vulnerable country among many languages and cultures. Boeke closes in what therefore seems to me to be a characteristically Dutch way, by pleading that his book's journey will help to make people better realize their place in the cosmic scheme of things, and in this way contribute to drawing the world closer together. Since I find his conclusion eloquent, I would like to close by quoting from it:

> When we thus think in cosmic terms, we realize that man, if he is to become really human, must combine in his being the greatest humility with the most careful and considerate use of the cosmic powers that are at his disposal.
>
> The problem, however, is that primitive man at first tends to use the power put in his hands for himself, instead of spending his energy and life for the good of the whole growing human family, which has to live together in the limited space of our planet. It therefore is a matter of life and death for the whole of mankind that we learn to live together, caring for one another regardless of birth or upbringing. No difference of nationality, of race, creed or conviction, age or sex may weaken our effort as human beings to live and work for the good of all.
>
> It is therefore an urgent need that we all, children and grown-ups alike, be educated in this spirit and toward this goal. Learning to live together in mutual respect and with the definite aim to further the happiness of all, without privilege for any, is a clear duty for mankind, and it is imperative that education be brought onto this plane.
>
> In this education the development of a cosmic view is an important and necessary element; and to develop such a wide, all-embracing view, the expedition we have made in these 'forty jumps through the universe' may help just a little. If so, let us hope that many will make it!

Post Scriptum.

By coincidence, in the same issue of *Scientific American* as this column appeared in, there was a short note in "Science and the Citizen" on the American nuclear arsenal. The information, compiled by the Center for Defense Information and the National Resources Defense Council, stated that the current stockpile amounted to some 30,000 nuclear weapons, 23,000 of which were operational. (An excellent way of visualizing this is shown in Figure 33-2, the last figure in the book.) The Reagan administration, it said, intended to build about 17,000 in the next ten years while destroying about 7,000, thus increasing the net arsenal by about 10,000 nuclear weapons.

This is roughly equivalent to ten tons of TNT per Russian capita. Now what does this really mean? Wolf H. Fahrenbach had the same nagging question, and he wrote to tell me what he discovered.

> Ten tons of TNT exceeds my numericity, so I asked a demolitions-expert friend of mine what one pound, ten pounds, 100 pounds, etc. of TNT could do. One pound of TNT in a car kills everybody within and leaves a fiery wreck; ten

> pounds totally demolishes the average suburban home; and 1,000 pounds packed inside an old German tank sent the turret to disappear in low overhead clouds. It could be reasonably suggested to the administration that most civilized nations are content with simply *killing* every last one of their enemies and that there is no compelling reason to have to ionize them.

Now this was interesting to me, because I happened to remember that the 241 marines killed in the recent truck-bombing in Beirut had been in a building brought down by what was estimated as one ton of TNT. Ten tons, if well placed, might have done in 2,400 people, I suppose. Ten tons is my allotment, and yours as well. That's the kind of inconceivable overkill we are dealing with in the nuclear age.

Another way of looking at it is this. There are about 25,000 megatons of nuclear weapons in the world. If we decode the "mega" into its meaning of "million", and "ton" into "2,000 pounds", we come up with 25,000 × 1,000,000 × 2,000 pounds of TNT-equivalent, which is 50,000,000,000,000 pounds to be distributed among us all, perhaps not equally—but surely there's enough to go around.

I find myself oscillating between preferring to see it spelled out that way with all the zeros, and leaving it as 25,000 megatons. What I have to remember is what "megaton" really means. Last summer I visited Paris and climbed the butte of Montmartre, from the top of which, at the foot of the Sacré Coeur, one has a beautiful view of all of Paris spread out below. I couldn't refrain from ruining my two friends' enjoyment of this splendid panorama, by saying, "Hmm... I bet one or two nicely placed megatons would take care of all this." And so saying, I could see exactly how it might look (provided I were a superbeing whose eyes could survive light and heat blasts far brighter than the sun). I know it seems ghoulish, yet it was also completely in keeping with my thoughts of the time.

Now if you just say to yourself "one megaton equals Paris's doom" (or some suitable equivalent), then I think that the phrase "25,000 megatons" will become as vivid as the long string of zeros—in fact, probably more vivid. It seems to me that this perfectly illustrates how the psychological phenomenon known as *chunking* is of great importance in dealing with otherwise incomprehensible magnitudes.

Chunking is the perception *as a whole* of an assembly of many parts. An excellent example is the difference between 100 pennies and the concept of one dollar. We would find it exceedingly hard to deal with the prices of cars and houses and computers if we always had to express them in pennies. A dollar has psychological reality, in that we usually do not break it down into its pieces. The concept is valuable for that very reason.

It seems to me a pity that the monetary chunking process stops at the dollar level. We have inches, feet, yards, miles. Why could we not have pennies, dollars, grands, megs, gigs? We might be better able to digest newspaper headlines if they were expressed in terms of such chunked units —provided that those units had come to mean something to us, as such. We all have a pretty good grasp of the notion of a grand. But what can a meg or a gig buy you these days? How many megs does it take to build a high school? How many gigs is the annual budget of your state?

Most numerically-oriented people, in order to answer these questions, will have to resort to calculation. They do not have such concepts at their mental fingertips. But in a numerate populace, everyone *should.* It should be a commonplace that a new high school equals about 20 megs, a state budget several gigs, and so on. These terms should not be thought of as *shorthand* for "million dollars" and "billion dollars" any more than "dollar" is a shorthand for "100 cents". They should be autonomous concepts—mental "nodes"—with information and associations dangling from them without any need for conversion to some other units or calculation of any sort.

If that kind of direct sense of certain big numbers were available, then we would have a much more concrete grasp on what otherwise are nearly hopeless abstractions. Perhaps it is in the vast bureaucracies' interest that their budgets remain opaque and impenetrable—but even that holds true only in the short run. Economic ruin and military suicide are not good for anybody in the long run—not even arms manufacturers! The more transparent the realities are, the better it is for any society in the long run.

* * *

This kind of total incomprehension extends even to the highest echelons of our society. Bucknell University President Dennis O'Brien recently wrote on the *New York Times* op-ed page: "My own university has just opened a multibillion-dollar computer center and prides itself that 90 percent of its graduates are computer-literate." And the Associated Press distributed an article that said that the U.S. federal debt ceiling had gone up to 1.143 trillion dollars, and then cited the latest figure for the debt itself as "$1,070,241,000". In that case, what's the hurry about raising the ceiling? These may have been typos, but even so, they betray our society's rampant innumeracy.

You may think I am being nitpicky, but when our populace is so boggled by large numbers that even many university-educated people listen to television broadcasts without an ounce of comprehension of the numbers involved, I think something has gone haywire somewhere. It is a combination of numbness, apathy, and a resistance to recognizing the need for new concepts.

One reader, a refugee from Poland, wrote to me, complaining that I had memorized hundreds of digits of π in my high school days without appreciating the society that afforded me this luxury. In East Block countries, he implied, I would never have felt free to do something so decadent. My feeling, though, is that memorizing π was for me no different from any other kind of exuberant play that adolescents in any country engage in. In a recent book by Stephen B. Smith, called *The Great Mental*

Calculators—a marvelously engaging book, by the way—one can read the fascinating life stories of people who were far better than I with figures. Many of them grew up in dismal circumstances, and numbers to them were like playmates, life-saving friends. For them, to memorize π would not be decadent; it would be a source of joy and meaning. Now I had read about some of these people as a teen-ager, and I admired, even envied, their abilities. My memorization of π was not an isolated stunt, but part of an overall campaign to become truly fluent with numbers, in imitation of calculating prodigies. Undoubtedly this helped lead me toward a deeper appreciation of numbers of all sizes, a better intuition, and in some intangible ways, a clearer vision of just what it is that the governments on this earth—West Block no less than East—are up to.

But there may be more direct routes to that goal. For example, I would suggest to interested readers that they attempt to build up their own numeracy in a very simple way. All they need to do is to get a sheet of paper and write down on it the numbers from 1 to 20. Then they should proceed to think a bit about some large numbers that seem of interest to them, and try to estimate them within one order of magnitude (or two, for the larger ones). By "estimate" here, I mean actually do a back-of-the-envelope (or mental) calculation, ignoring all but factors of ten. Then they should attach the idea to the computed number. Here are some samples of large numbers:

* What's the gross state product of California?
* How many people die per day on the earth?
* How many traffic lights are there in New York City?
* How many Chinese restaurants are there in the U.S.?
* How many passenger-miles are flown each day in the U.S.?
* How many volumes are there in the Library of Congress?
* How many notes are played in the full career of a concert pianist?
* How many square miles are there in the U.S.? How many of them have you been in?
* How many syllables have been uttered by humans since 1400 A.D.?
* How many "300" games are bowled in the U.S. per year?
* How many stitches are there in a stocking?
* How many characters does one need to know to read a Chinese newspaper?
* How many sperms are there per ejaculate?
* How many condors remain in the U.S.?
* How many moving parts are in the Columbia space shuttle?
* How many people in the U.S. are called "Michael Jackson"? "Naomi Hunt"?
* What volume of oil is removed from the earth each year?
* How many barrels of oil are left in the world?
* How much carbon monoxide enters the atmosphere each year in auto exhaust fumes?

* How many meaningful, grammatical, ten-word sentences are there in English?
* How long did it take the 200-inch mirror of the Palomar telescope to cool down?
* What angle does the earth's orbit subtend, as seen from Sirius?
* What angle does the Andromeda galaxy subtend, as seen from earth?
* How many heartbeats does a typical creature live?
* How many insects (of how many species) are now alive?
* How many giraffes are now alive? Tigers? Ostriches? Horseshoe crabs? Jellyfish?
* What are the pressure and temperature at the bottom of the ocean?
* How many tons of garbage does New York City put out each week?
* How many letters did Oscar Wilde write in his lifetime?
* How many typefaces have been designed for the Latin alphabet?
* How fast do meteorites move through the atmosphere?
* How many digits are in 720 factorial?
* How much is a brick of gold worth?
* How many gold bricks are there in Fort Knox? How much is it worth?
* How fast do your wisdom teeth grow (in miles per hour, say)?
* How fast does your hair grow (again in miles per hour)?
* How fast is Venice sinking?
* How far is a million feet? A billion inches?
* What is the weight of the Empire State Building? Of Hoover Dam? Of a fully loaded jumbo jet?
* How many commercial airline takeoffs occur each year in the world?

These or similar questions will do. The main thing is to attach some concreteness to those numbers from 1 to 20, seen as exponents. They are like dates in history. At first, a date like "1685" may be utterly meaningless to you, but if you love music and find out that Bach was born that year, all of a sudden it sticks. Likewise with this secondary meaning for small numbers. I can't guarantee it will work miracles, but you may increase your own numeracy and you may also help to increase others'. Merry numbers!