

17600	44165	13288	21897	29692	09820	47937	28125	74363	14070	75332
17601	64553	71275	59710	96838	19125	47805	86569	66843	13992	71758
17602	08734	68919	58268	81797	39081	54855	40790	78442	45637	24565
17603	34689	82532	05640	81699	42644	58136	82477	79851	42954	42565
17604	96381	96387	77323	98977	33246	84059	96505	21419	96017	93464
17605	68099	57848	66633	63172	25993	66390	68963	52859	45475	33291
17606	10388	45585	39883	11136	13772	84701	13882	01641	82485	05557
17607	17979	42128	64164	22200	39987	16598	61654	06667	77928	47179
17608	45852	79716	94810	95520	44793	46386	36534	16964	78772	13403
17609	98015	66227	14076	30854	73752	11801	84730	11571	04314	39338
17610	13814	41470	98284	04408	65093	57907	51851	12720	18064	96515
17611	72519	78719	66131	23993	74347	79042	22600	00962	70595	46516
17612	03536	76505	10481	40172	94339	53624	31745	04222	54301	80363
17613	86248	73520	04485	41459	85352	49384	51885	10852	65322	65472
17614	66842	86998	44532	38597	87423	88063	73803	59590	10174	86191
17615	04902	32957	44799	90914	25572	79017	78887	94096	85225	21153
17616	45478	73392	30745	54492	41232	66093	55241	16232	27482	86385
17617	58119	92219	14214	50332	93952	22506	75434	28771	33557	25884
17618	84663	05917	48868	02408	91187	11031	59939	97149	18336	37394
17619	15006	72976	73334	59990	80824	45871	27081	44625	24194	39784
17620	37154	45999	93071	25965	96363	61747	66505	13178	34649	67397
17621	37427	70943	79642	45355	14855	55282	45411	01773	08702	11962
17622	75403	90631	23321	48939	71599	62160	24538	06654	13875	65120
17623	45154	50924	92034	59766	42005	98036	22254	61897	08681	49281
17624	91085	20839	46831	00884	73126	71163	24302	48886	50017	85596
17625	78271	40424	87355	99928	76920	30663	46996	18023	76435	51925
17626	54029	32780	93160	57546	77945	31410	20481	44858	64476	46941
17627	23534	35412	63871	19117	55346	16228	84433	98960	89044	44805
17628	92767	70032	79644	93751	79177	56863	02712	03906	41685	04183
17629	33505	69043	44839	37250	53233	99293	79440	04677	71153	12828
17630	92103	51231	77397	75191	05741	46489	78309	19861	57819	88807
17631	65835	71066	02338	23531	39970	59453	37252	11789	78467	37814
17632	77711	68170	11209	75619	70738	78470	40312	14522	92626	96833
17633	70989	10117	87486	98240	88607	65130	65761	60848	90999	30210
17634	99240	55016	37052	06431	12364	86658	50516	65705	58460	71690
17635	39857	88095	50544	93912	62074	84982	13564	36570	31292	15136
17636	99839	25488	74561	29858	50876	11241	49809	82350	59760	56664
17637	54935	88698	65322	12870	80740	43392	35426	55905	12041	89327
17638	55062	35696	48490	69207	09511	58423	72175	31289	35347	74299
17639	88384	61009	50998	30826	53295	27499	33753	23301	96736	63624
17640	92744	06115	67100	80515	42861	15094	30524	95191	52564	49676
17641	28932	36793	91356	60265	91059	36261	97082	02698	37843	49068
17642	40573	33598	78476	35301	14875	38035	65064	20732	50658	93901
17643	51926	24751	13834	21960	35006	16166	94046	65237	77965	78338
17644	73663	75788	91024	38226	25378	52296	63064	50616	45521	13992
17645	18767	94575	39310	08046	86327	61028	79876	96211	93553	95586
17646	61484	88255	84328	13377	74997	59692	08937	69909	69175	31400
17647	86565	11477	03819	68703	09578	02389	23004	33566	60312	79310
17648	14624	72368	82187	61133	50419	50136	07609	84623	80483	10453
17649	81678	22258	26592	77602	55355	16003	50591	15294	46367	11105

The first two pages of this bulletin originally appeared under the title “A Perfectly Normal Distribution of Precisely No Information” in *Nozone X: Forecast* (New York: Nozone, 2008), and also in *A Couple Thousand Short Films About Glenn Gould* (London: Film and Video Umbrella, 2008). The original text has been considerably appended here, and egregious errors retroactively addressed.

Cover image: p.353 from *A Million Random Digits (With 100,000 Normal Deviates)*, published by RAND Corporation (1955)

Just after World War II, The RAND Corporation was quietly working on a massive book of numbers. *A Million Random Digits (With 100,000 Normal Deviates)* was published by The Free Press in 1955 after almost ten years of meticulous production. The volume is comprised of page after page of numbers—mathematical tables filled with random digits. (A typical page (picked at random) from the 1966 printing is reproduced on the cover of this bulletin.) The random number bible has passed through three editions, multiple printings, and is currently available as both a soft format paperback book and as a text data file downloadable directly from RAND.

A Million Random Digits was produced out of an increasing demand at RAND from the onset of the Cold War. Random numbers are necessary for all kinds of experimental probability procedures including game simulations and scenarios, weather forecasting, message encryption and compression, financial market projections and any complex statistical model that attempts to predict future behavior based on past observation. In order to reproduce probabilistic situations, purely random numbers are critical as numerical starting points and/or additional data sources. Without them, these statistical projections, or Monte Carlo models as they are commonly called, will show biases based on their starting conditions that make their forecasts (mathematically) useless.

Producing a random digit is complex. To make the tables in *A Million Random Digits*, RAND engineers created an electronic roulette wheel with 32 possible values by measuring the decay of a radioactive molecule gated by a constant frequency pulse. These regular electric signals (either on or off, 100,000 times a second for 10 seconds) were run through a five digit binary counter to produce a 5-bit number with 32 possible values. The binary number was converted to decimal and only the final digit was retained to create the 1,000,000 random digits. These values were fed into an IBM punch machine to produce 20,000 computer punch cards with 50 digits each. (Punch cards were then the only practical way to both store and input information into a digital computer.) However, when analyzing this first attempt, RAND engineers detected a bias. Employing a standard statistical goodness-of-fit test to measure the data's conformity to a bell-shaped or "normal" curve, the sampled numbers did not match closely enough to the normal distribution of values which would indicate purely random digits. Each number was added modulo 10 (divide by 10 and

use only the remainder) to the corresponding digit on the previous card to yield a new set of random values with an almost perfectly normal distribution. Random digit tables were then printed on an IBM 856 Cardatype and reproduced as pages for the book. Proofreading was redundant given the nature of the information. Nonetheless, every twentieth page was proofed and every fortieth was summed against the original punch cards.

Using a random digit from the book is not much simpler. Instructions are included in the introduction and read as some kind of cabalistic incantation:

Open the book to an unselected page of the digit table and blindly choose a five-digit number; this number with the first digit reduced modulo 2 determines the starting line; the two digits to the right of the initially selected five-digit number are reduced modulo 50 to determine the starting column.

While RAND engineers were producing *A Million Random Digits*, two American mathematicians, MIT professor Norbert Wiener and Bell Labs researcher Claude Shannon were simultaneously creating rigorous mathematical models of communication processes which together are known as Information Theory. In this widely-applied framework, information is defined as the amount that one value can tell you about the next value. For example, the value 12:32 PM tells you that the next should be 12:33 PM—therefore, 12:32 PM has a high Information content. Or, a temperature of 71 ° F gives you a pretty good idea that the next value will remain in a limited range not far from 71 ° F. The fantastic achievement of this book is that each random digit in it tells you precisely nothing about the next. (This is the point.) *A Million Random Digits* is then a book that contains exactly, rigorously, meticulously and absolutely *no* information.

Four years have passed since the last paragraph.

In the meantime, it's become clear that I got it perfectly wrong when the previous two pages of this text were first published under the title, "A Perfectly Normal Distribution of Precisely No Information." *A Million Random Digits ...* does not contain *no* information, but, exactly

the opposite. The digits that construct the body of this volume contain nothing but information.

“I” IS FOR INFORMATION

Claude Shannon’s 1948 paper “A Mathematical Theory of Communication,” published in volume 27 of *The Bell System Technical Journal*, introduced the technically proper term “information” and outlined most of the not-yet-field of Information Theory. (The paper is so often referenced that academic journals provide a shortcut notation— just drop in “\shannon48” and the complete bibliographic details are automatically appended.) I’d read the paper before, but obviously missed its argument or at least the proper use of its terminology. In it, Shannon reloaded the common English word “information” with a new, technically precise and mathematically operative definition. He did the same with “entropy.” The new definitions are bound up in subtle mathematical relationships.

The fundamental problem of communication, according to Shannon, is that of “reproducing at one point either exactly or approximately a message selected at another point.” The message communicated is one selected from a set of possible messages. “Information,” for Shannon, was then the *substance* transferred by any communication. It is the material that gets moved from here to there. “A Mathematical Theory of Communication” proceeded from this premise to quantify and precisely model a generic communication system. As Shannon was an electronic engineer, much of the argument is carried in its mathematics—so in going back to the source of my “information” confusion, I decided this time to re-read the paper and work my way slowly through its equations.

The first was simple enough and essential, relating the amount of information communicated to the probability of picking any one particular message from a group of possible messages. For example, answering “yes” or “no” is one choice from a set of two possible messages. Messages can be of any length or complexity, and Shannon emphasized that the meaning of the message was irrelevant to the technical problem he was addressing. He was interested only in its transmission, a matter of reproduction, not interpretation.

Another way to describe the amount of information would be to specify the *freedom of choice* in picking any one from a set of possibles. For example, if I were to ask you “Do you like tea?” then I could be reasonably certain that you will reply either “Yes” or “No.” But if instead, I enquired as to what kind of tea you like, your selected answer would come from the near-infinite set of tea varieties. (Today I’d answer, “Simple, Assam.” Ask me again tomorrow.) The information content of the first answer is relatively low, I know it will be either “Yes” or “No.” The second answer’s information is considerably higher, as the range of possible answers is much greater. Any one particular choice carries more information, more consequence—it matters more. The less certain you are of my answer, the more weight my reply carries. Simply, the amount of information measures the change in your uncertainty produced by my answer to the question.

The mathematical relation that describes the amount of information (I) when one message is selected from a set of possible messages is:

$$I = \log_2 N$$

where N is the total number of possible messages. Extracting the logarithm (in base 2, rather than the base 10 of our usual decimal arithmetic) is only a matter of asking what exponent of 2 will equal the number in question. For example, 2 to the what-th power equals 8? The answer is 3, so that $\log_2 8 = 3$. Anyway, we can then apply this formula to the 1,000,000 numerals from the RAND Corporation’s book to find the total amount of information contained within. (We will here assume that since each number is meticulously “random,” then the selection of any one number is discrete, or has no effect on the selection of the next digit.) Any random digit is selected from a set of ten possible choices (0–9), so to find the amount of information contained in any one digit, we let $N = 10$ and

$$I = \log_2 10$$

$$I = 3.32 \text{ bits}$$

Because each digit in the book is completely independent of every other digit, then we can just multiply that amount by the number of digits to

get the total information contained in the book, which you'll notice is considerably more than the "no" information I originally suggested:

$$I = 3.32 \text{ bits} \times 1,000,000 \text{ digits}$$

$$I = 3,320,000 \text{ bits}$$

Try this same exercise with $N = 2$ (two possible digits, 0 or 1) and you'll find that $I = 1,000,000$ bits or 1 bit for each digit.

Our answers for I are given in "bits," a compressed neologism Shannon also proposed in this paper: "the resulting units may be called binary digits, or more briefly, bits." A bit represents a choice between two possibilities. It may be "on" or "off," "black" or "white," "0" or "1," "A" or "B." The bit is information's atom, the smallest indivisible unit, its essential measure, or as anthropologist Gregory Bateson described a bit some years later, it's *the difference that makes a difference.*

As soon as information could be quantified, measured and relayed in consistently measured chunks as bits, then it no longer mattered what kind of information was being relayed, what it meant, or to whom. Information was freed from meaning and now became a thing, as real as water and at least as fluid. It could be carried in the words on the pages of a book, by a secret whispered in confidence, through currents crackling over telegraph wires, and most consequently via electrical charges pulsing through the silicon valleys of a computer chip.

"H" IS FOR ENTROPY

Claude Shannon had another letter as well, which followed the I of Information and was integrally tied to it— H for "entropy." He borrowed the term from physics where the Second Law of Thermodynamics describes entropy as the inevitable one-way tendency of any system to fall into disorder. Mathematician John von Neumann, who had already done extensive work in expanding classical entropy into the fuzzy maths of quantum mechanics, pointed Shannon to the word as a proper name for his concept of informational uncertainty, suggesting appropriately enough:

... no one really knows what entropy really is, so in a debate you will always have the advantage.

Like “information,” Shannon overloaded “entropy” with a new, precisely technical definition—it is the measure of uncertainty in the value of a random variable. So the entropy in the outcome of a perfect coin toss is a maximum value, or 1 as it is equally likely that the toss is either heads or tails. If the coin is weighted towards heads then the entropy decreases—we can now guess that it is a little more likely to be heads than tails, so our uncertainty has been reduced. If the coin has two heads, then we can be rather certain that any coin toss will produce heads. Therefore, our uncertainty is reduced to nothing (or entropy (H) = 0).

The general (now canonical) equation for entropy that Claude Shannon published in his 1948 paper has the form:

$$H = - \sum_{i=1}^N p(i) \log_2 p(i)$$

I shouldn't go into too much detail regarding this mathematical string of para-alphabetical symbols, but there a couple of things you should know: (1) the large, bent Σ is a Sigma and means the sum over a range of values (from 1 to N here); and (2) the italic p stands for probability (in this case of picking the value i). It is enough for now to understand that the entropy in a certain quantity of symbols (say in the digits of the RAND book) is measured by summing the information content of each weighted by their relative probabilities. As the individual choices become equally likely (or random) then the entropy works its way towards a maximum value, which is equivalent to the brute information content of the message. Any reduction in uncertainty, or randomness, reduces the entropy.

So, for a purely random digit between 0–9, what's its H ?

Let $N = 10$ and

$$H = - \sum_{i=1}^N p(i) \log_2 p(i)$$

$$H = - \sum_{i=1}^N .10 \log_2 .10$$

$$H = - \sum_{i=1}^N .10 \times -3.32$$

$$H = \sum_{i=1}^N .332$$

$$H = 3.32 \text{ bits}$$

Looks familiar, right?

If you did the same gymnastics with a less purely random digit between 0 and 9—let's say that picking a 7 is six times more likely than picking any other digit—then the resulting entropy is lower, $H = 2.87$ bits. You could now guess that it is more likely to pull a 7 than any other digit, and so the situation is a little bit less uncertain; hence, H goes down.

But RAND engineers wanted the biggest H possible, and worked ruthlessly to produce a series of digits which contain absolutely maximum entropy. So if *A Million Random Digits ...* contains perfect disorder (each number is fastidiously, precisely random), and it also contains total information, then it seems that there must be some essential relationship between information and randomness.

Information is not the same thing as randomness, but instead, they are complements—tied together in a push-me-pull-you arrangement. One relies on the other. It turns out that pure randomness is fundamental to the digital communications of our so-called information age, and in a circular snake-eating-its-own-tail ouroboros conundrum, equally impossible to produce within the deterministic logic of a binary electronic computer.

“ π ” IS FOR PI

Concluding his internal RAND Corporation report of June 1949 on the extended and difficult process of developing the million random digits, George W. Brown says:

My own personal hope for the future is that we won't have to build any more random digit generators. It was an interesting experiment, it fulfilled a useful purpose, and one can do it again that way, if necessary, but it may not be asking too much to hope that this addition property, perhaps, or some other numerical process, will permit us to compute our random numbers as we need them.

This dream of programmatically (computationally) producing a purely random string of digits as Brown wished remains unresolved. As in the book, the definition of a random string of numbers requires that it contains no pattern—each number is completely discrete of any other number. There is no way to reduce the sequence of digits to any formula, or any program. For example, the first million digits of π appear at first glance to be random enough:

$$\pi = 3.14159265 \dots 779458151$$

but that same number can be easily reproduced by the compact formula (or program) of dividing the circumference of a circle (length around the edge) by its diameter (length from one side to the other.) The recipe is:

$$\pi = C / d$$

So then the digits in π are not random at all. They are precisely accounted for by a simple, short program. (Recall that a random sequence cannot be compressed—the only way to determine the next digit is to randomly pick the next digit.) So a truly random number cannot be defined by a program, or definite method. Binary electronic computers are finite state machines, designed to be entirely predictable and run only by a set of instructions written as a computer program. Therefore, by definition, a computer cannot produce a random number as a random number cannot be computed.

The best a digital computer can do is to produce what is called a PSEUDO- random number. There are many increasingly sophisticated algorithms employed whose quality of randomness (or entropy) is high enough for a majority of applications, but these are inevitably compromised by the fundamental impossibility of their task.

The only way to produce a purely random number in a computer is to join it to the quantumly messy world of life outside of its box. This is precisely the strategy taken by a number of competing projects that generate and release random numbers via the World Wide Web. www.random.org offers a true random number generator by using three radios tuned in-between stations to capture atmospheric noise as a data source (or seed) for its continuous generation of noisy digits. On March 14 2012 at 1:32 PM, I requested and received the following random number between 1 and 1,000,000:

566662

which doesn't look so random to me, but that seems to be the trick with real randomness—our own desires to produce patterns filters our empirical experience. Since going live in October 1998, random.org has offered up 1,094,739,583,783 bits of pure entropy and counting. Competing random number sites include Hotbits (www.fourmilab.ch/hotbits) which employs the radioactive decay of the Cesium-137 nucleus (you can hear the bits as they are made) and even a setup that harnesses the random activity of a lava lamp to create the seed of a random number. www.lavarand.org takes continuous digital snapshots of the internally chaotic states of a lava lamp at a given moment and uses this collection of bits passed through a hash function to seed a high-octane pseudorandom generator called the Blum Blum Shub.

The irony of expiring so much effort to produce something that surrounds us everywhere we look is not lost on Random.org's founder, Mads Haahr. In a *New York Times* interview from 2001, he acknowledged:

It was a bit like selling sand in the desert. But it's not quite like that, because the noise you're getting from Random.org is pure in a way; it's different from the hustling bustling cacophony of the information age.

Producing anything that's pure, even noise, takes effort.

So why does it matter?

When RAND released its book, the numbers were most significantly valuable for statistical and experimental models and mathematical projections. Since then and with the rise of computer networks, they have become indispensable—random numbers are essential to using computers to communicate, securely, robustly, and consistently. Purely random digits are the degree zero of cryptography (“secret writing”), producing the unbreakable foundation of a secure encoding system. These are used in, for example, online payment systems and the safeguarding of sensitive databases. More fundamentally, they are employed to securely identify a particular computer on a network or ensure that a message addressed to one person reaches that one person and nobody else.

The certainty that users have in the reliability of a communications system is likely to affect what gets said through it. If the mail ran occasionally and a letter only sometimes reached the intended address, I probably wouldn't say something important to you in a letter. A secret whispered in confidence relies on the trust between two communicating parties: “Shhhhhh ... do you want to maybe step out for a drink?”

One widely employed system for encoding digital communication today is called Public-key cryptography and was properly introduced for use with digital computers in 1976, by Whitfield (“Whit”) Diffie and Martin Hellman. The Diffie-Hellman arrangement requires the production of two separate keys for use in transmitting a secure message: one private key and one public key. By publishing the public key for use by anyone wishing to communicate with its owner and keeping the other secret, the message is securely passed. A real-world analogy would be something like: Anne has a public mailbox with a slot. If Bill would like to get a message to Anne, he drops it in. Only Anne, with her mailbox key can recover the message.

The corresponding electronic keys are actually very large numbers produced mathematically from an also-very-large random seed. The public key is produced by multiplying two long prime numbers—undoing this operation into its component factors is prohibitively difficult. By adding

the random seed, the process becomes impossible to duplicate. (Factoring a giant number into prime numbers is a bit like trying to un-mix two colors of paint in a can.)

And yet all of this depends essentially on the randomness of the initial number for its strength. Early this year a team of European and American mathematicians and cryptographers uncovered a significant flaw in a currently very-widely used data encryption algorithm and published it in the whimsically-titled paper “Ron Was Wrong, Whit Was Right” (a sly reference to Whit Diffie and Ron Rivest, two key players in this arena). The researchers examined 7.1 million public keys and discovered that, based on the weakness (impurity) of their random number generation, 27,000 of these were immediately vulnerable to simple cracking. Cryptographic weaknesses stemming from insufficiently random numbers have made the news before. In 1995, two researchers at University of California Berkeley discovered a flaw in the Netscape browser, and just last year, a similarly lazy approach in the software security of Sony Playstation 3 led to a massive breach of personal data for over 75 million people.

But it isn't just the security of personal information or the passing of secret messages that's at stake. Claude Shannon had mapped this territory already—recall him describing:

The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point.

He goes on to detail how the transmission of any message through a channel from one point to another **requires** encoding (encrypting) of that message. No communication is possible without this step.

As communication is ever-increasingly electronic and digital, then it is exactly this encryption that ensures that a message (its information) makes the journey confidently, predictably, and securely. And these encoding schemes and transfer protocols, for example the two-headed public/private key encryption scheme just described, rest firmly on the quality of the random numbers that stamp their exchanges.

We use machines to communicate from person-to-person. We email, we chat, we post, we search; and each time the transaction relies on one computer speaking to another. If the elaborate mathematical dance of these encryption protocols is the language that allows machine-to-machine communication, then the pure entropy of *A Million Random Digits ...* is its alphabet.

*