Computation and Civilization

Charles H. Bennett,
IBM Research Yorktown
HPC-UF Boston
Sept 10, 2013

X-ray of the Antikythera Mechanism, an early computer, Athens

The irony of inscrutability—having information but not being able to take advantage of it, because of limited computing abilities.

• What was the world like before HPC?

• What would it be like if all computations were easy?

• What fundamental limitations are there on our ability to predict the future and remember the past?

• Can computation theory be used to arrive at a non-anthropocentric criterion of ethical or cultural value, of that which civilization ought to foster and preserve.

Prehistory: Before HPC

- Accounting, Counting, Writing 4000-3000 BC
- Abacus, Place notation, Zero
- Al-Khwārizmī ca 825
- Logarithms 1614, Slide rule
- Mechanical calculators 19th-20th C.

30°

′	log sin 9–10	log tan 9–10	log cot 0	log cos 9–10	′
0	.69 897	.76 144	.23 856	.93 753	60
1	919	173	827	746	59
2	941	202	798	738	58
3	963	231	769	731	57
4	984	261	739	724	56

Early 20th Century tables of base-10 logarithms and logs of trig functions

Bulla and tokens from Susa, Mesopotomia, 4th Millennium BC (Louvre)

In his excellent book *The Information,* James Gleick recounts how, before the discovery of information theory, many of the important functions of the Internet were performed by other ingenious means, more slowly and expensively.

Beacon fires relayed news of the fall of Troy 12th century BC

Talking drums of West Africa reproduced pitches and prosody of local tonal languages over long distance, using stereotyped verbosity to compensate for lost consonants and vowels.

Come back home becomes (in English translation)
Make your feet come back the way they went,
Make your legs come back the way they went,
Plant your feet and your legs below,
In the village which belongs to us.

A similar phenomenon has occurred with **Whistled Languages**, indigenous to several mountainous regions around the world.

During the 20th century, computers became powerful enough to do things that are quite useful and in many cases would have looked like magic100 years ago:

- gravitational slingshots
- large optimization problems
- computational physics and chemistry
- hydrodynamics
- cryptanalysis

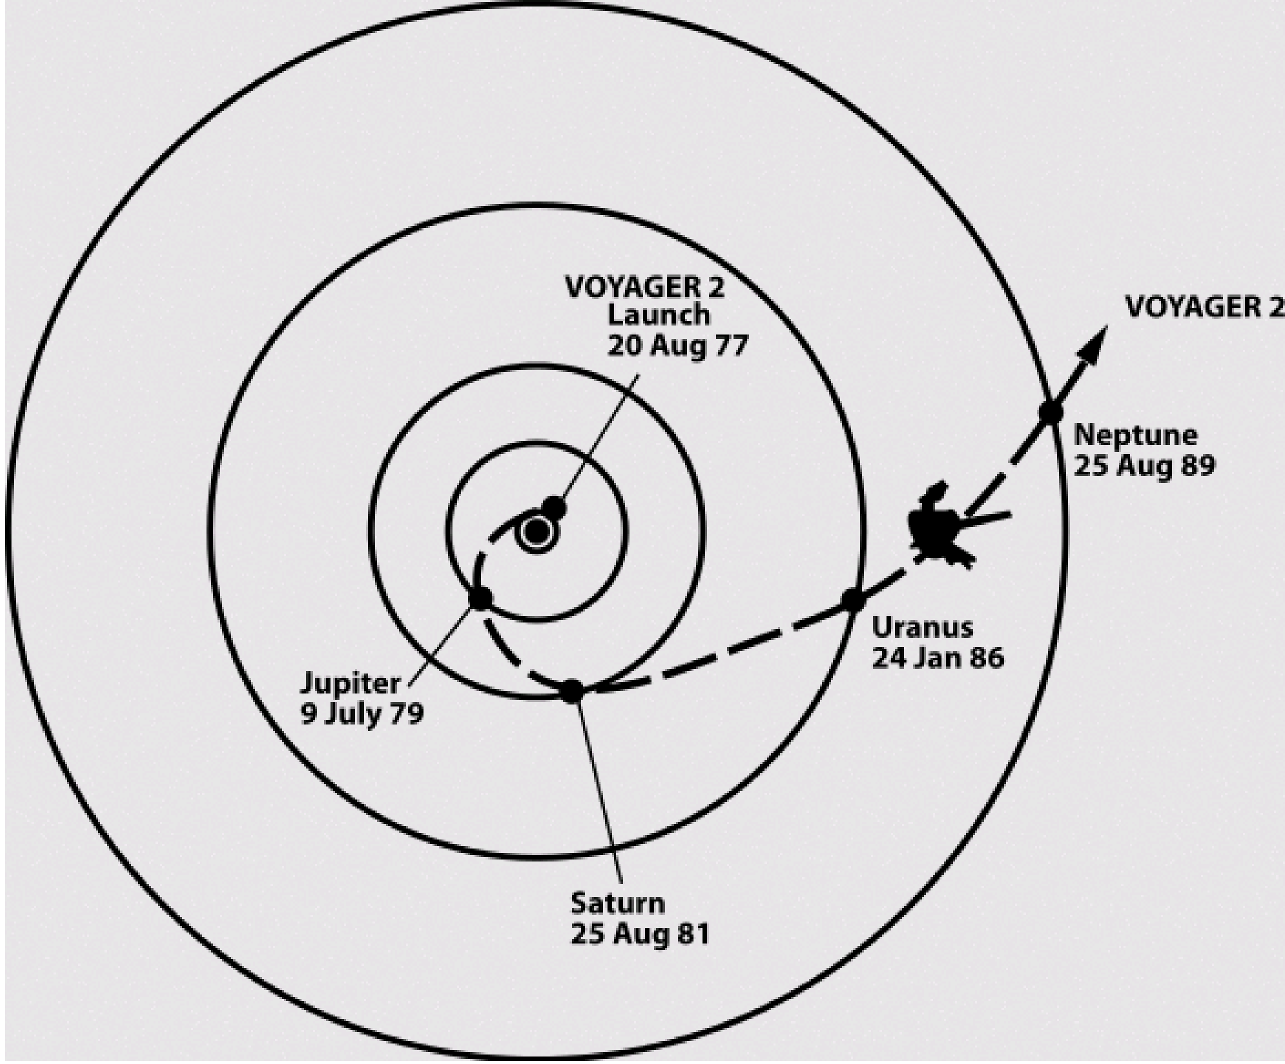

How far can we expect this to continue? Would a world in which no computation was too hard look totally different and magical?

What would it mean for no computation to be hard?
Before attempting to answer that, consider this model of "all computations" devised by Gregory Chaitin in 1975. It is a modern form of the old idea of a monkey at a typewriter eventually typing the works of Shakespeare. Of course a modern monkey uses a computer instead of a typewriter.

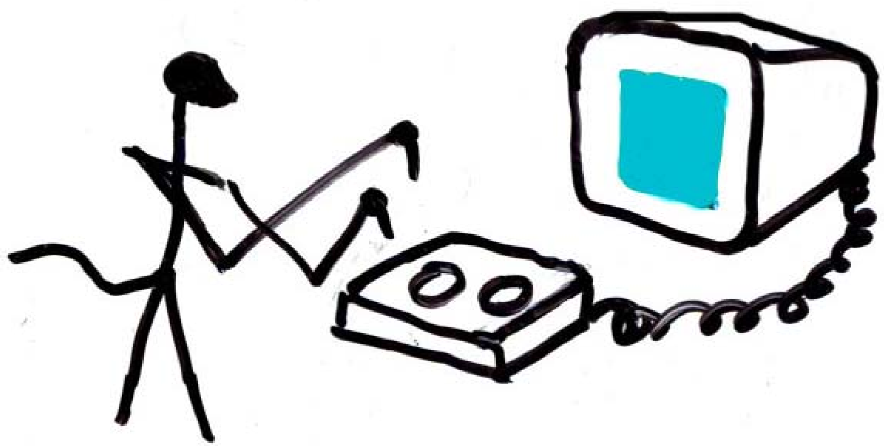

A monkey randomly typing 0s and 1s into a universal binary computer has some chance of getting it to do any computation, produce any output .

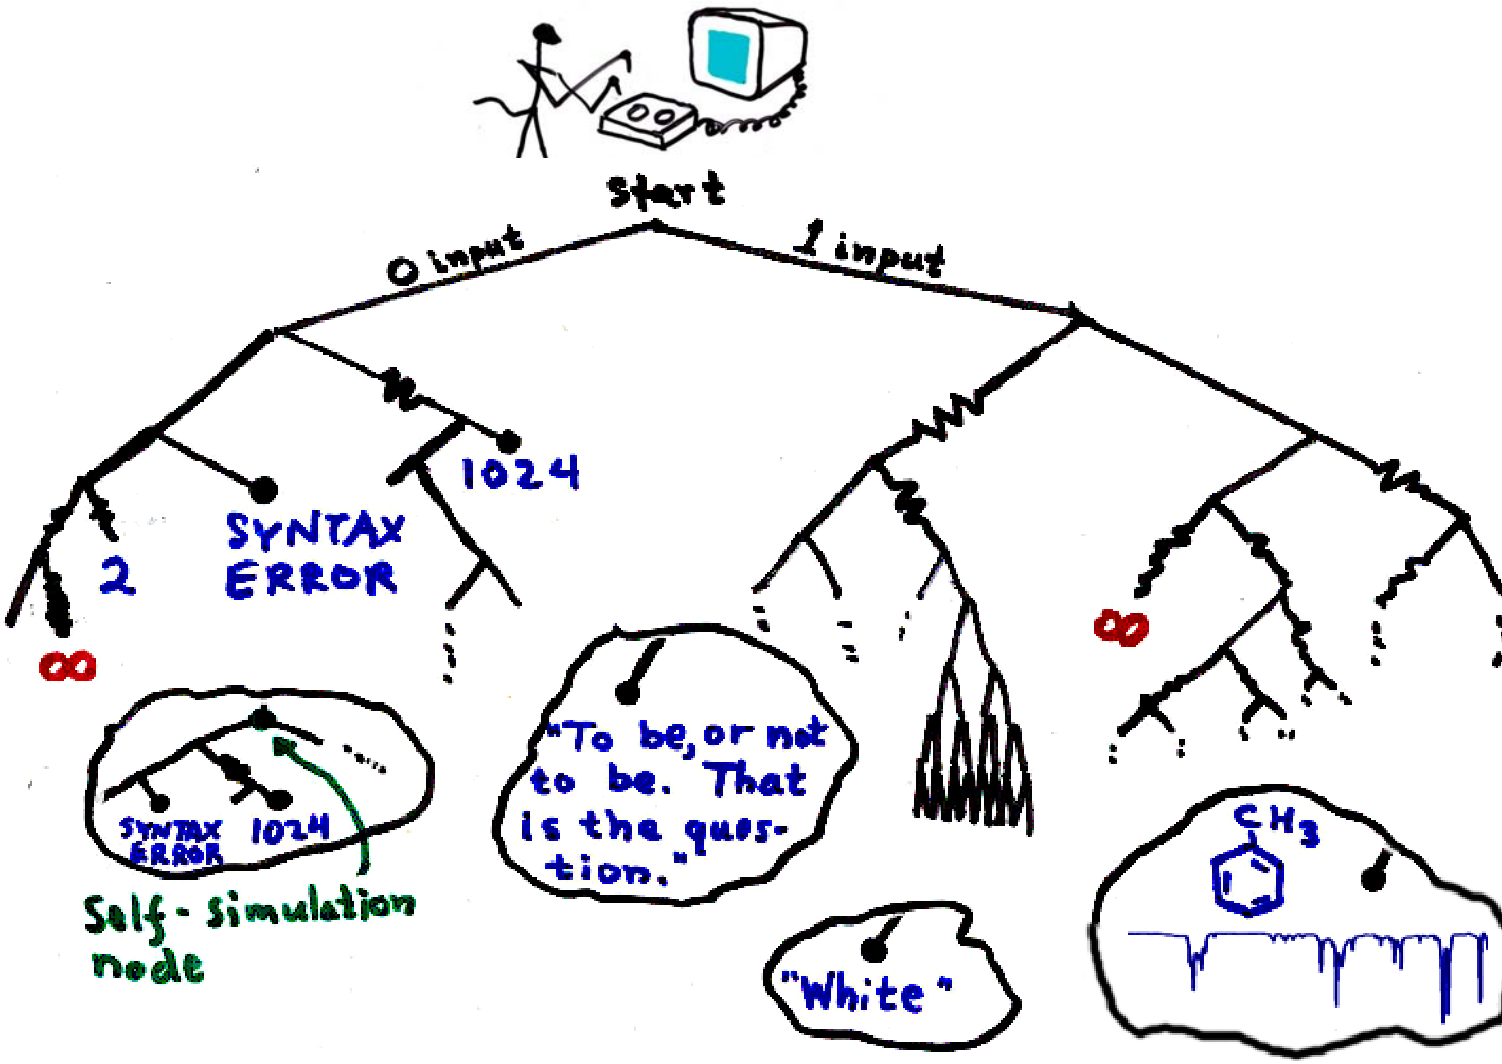

The input/output graph of this or any other universal computer is a microcosm of all cause/effect relations that can be demonstrated by deductive reasoning or numerical simulation.

The notion of "unlimited computing power" is somewhat problematic in light of Turing's demonstration of the unsolvability of the halting problem.

Turing went on to consider a hypercomputer consisting of an ordinary universal computer equipped with an **oracle *K*** for the halting problem. He showed that a ***K***-equipped computer still cannot solve some problems, in particular ***its own*** halting problem. Nevertheless it would be very powerful, answering almost all problems of practical interest as quickly as they could be asked.

A hypercomputer could not find deterministic solutions for problems whose solutions, owing to quantum and chaotic effects, were intrinsically probabilistic, but it could quickly calculate and optimally sample the relevant probability distributions.

Thus even with such a magical resource, some HPC applications, such as long range weather forecasting, would be limited by intrinsic probabilism, as well as by the quality and quantity of initial data.

There is no reason to suppose that anything remotely like a hypercomputer can exist in the physical world, so it is more realistic to consider other models of enhanced computation.

• A **quantum computer.** Consensus is that there's no obstacle in principle to its eventual construction, but many practical obstacles. Significant speedup for some problems over a classical computer, no speedup for others.

• An unexpected discovery in theoretical computer science such as **P=NP** or **P=PSPACE.** Such a discovery would not make computers stronger, but would imply that many problems suspected of being difficult, e.g. protein folding, are exponentially easier than formerly thought, being doable on an ordinary Turing machine in polynomial time (**P**). It is hard to imagine this kind of world, and most computer scientists avoid thinking about it.

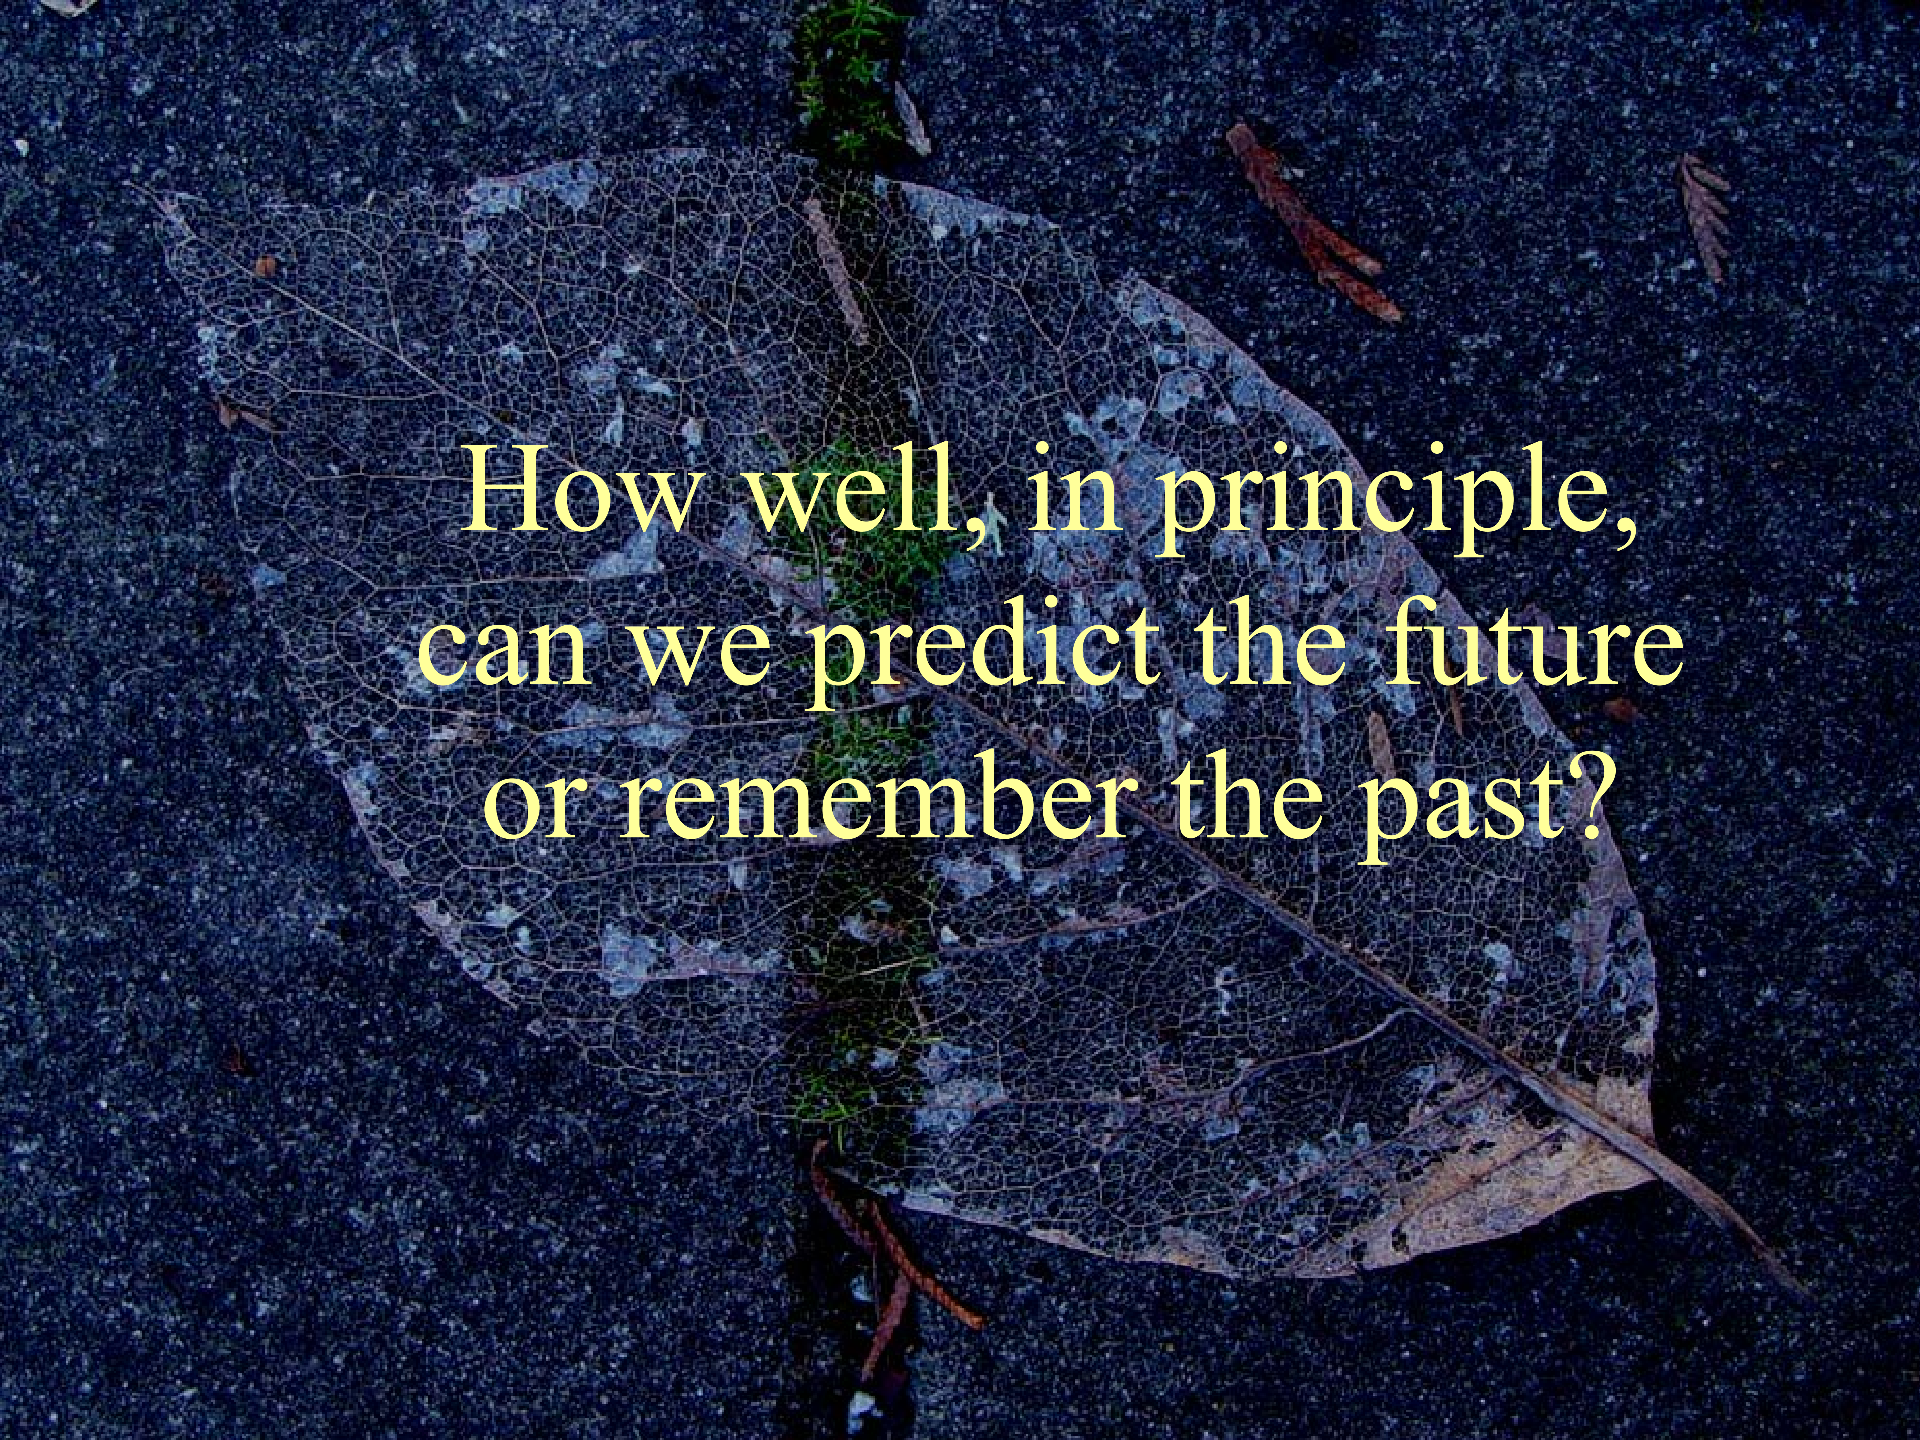
How well, in principle,
can we predict the future
or remember the past?

Reasoning from classical mechanics, Laplace thought the future and past were fully determined by the present, but attributed the perceived ambiguity of the future to our imperfect knowledge of the present, and/or our lack of sufficient computing power to calculate the future. An omniscient God would know past, present, and future.

Quantumly, the future is less determined than Laplace imagined. Even an omniscient God would not be able to predict whether a particular radioactive atom will decay within its half life.

In our macroscopic world, we remember the past much better than we can predict the future. One can now scan all the books in Google Books to see how the frequency of various phrases have varied over time. The phrase "1970" is mentioned rarely before that year, often immediately after, then with declining frequency.

Google books Ngram Viewer

Year of publication:

1960 1970 1980 1990

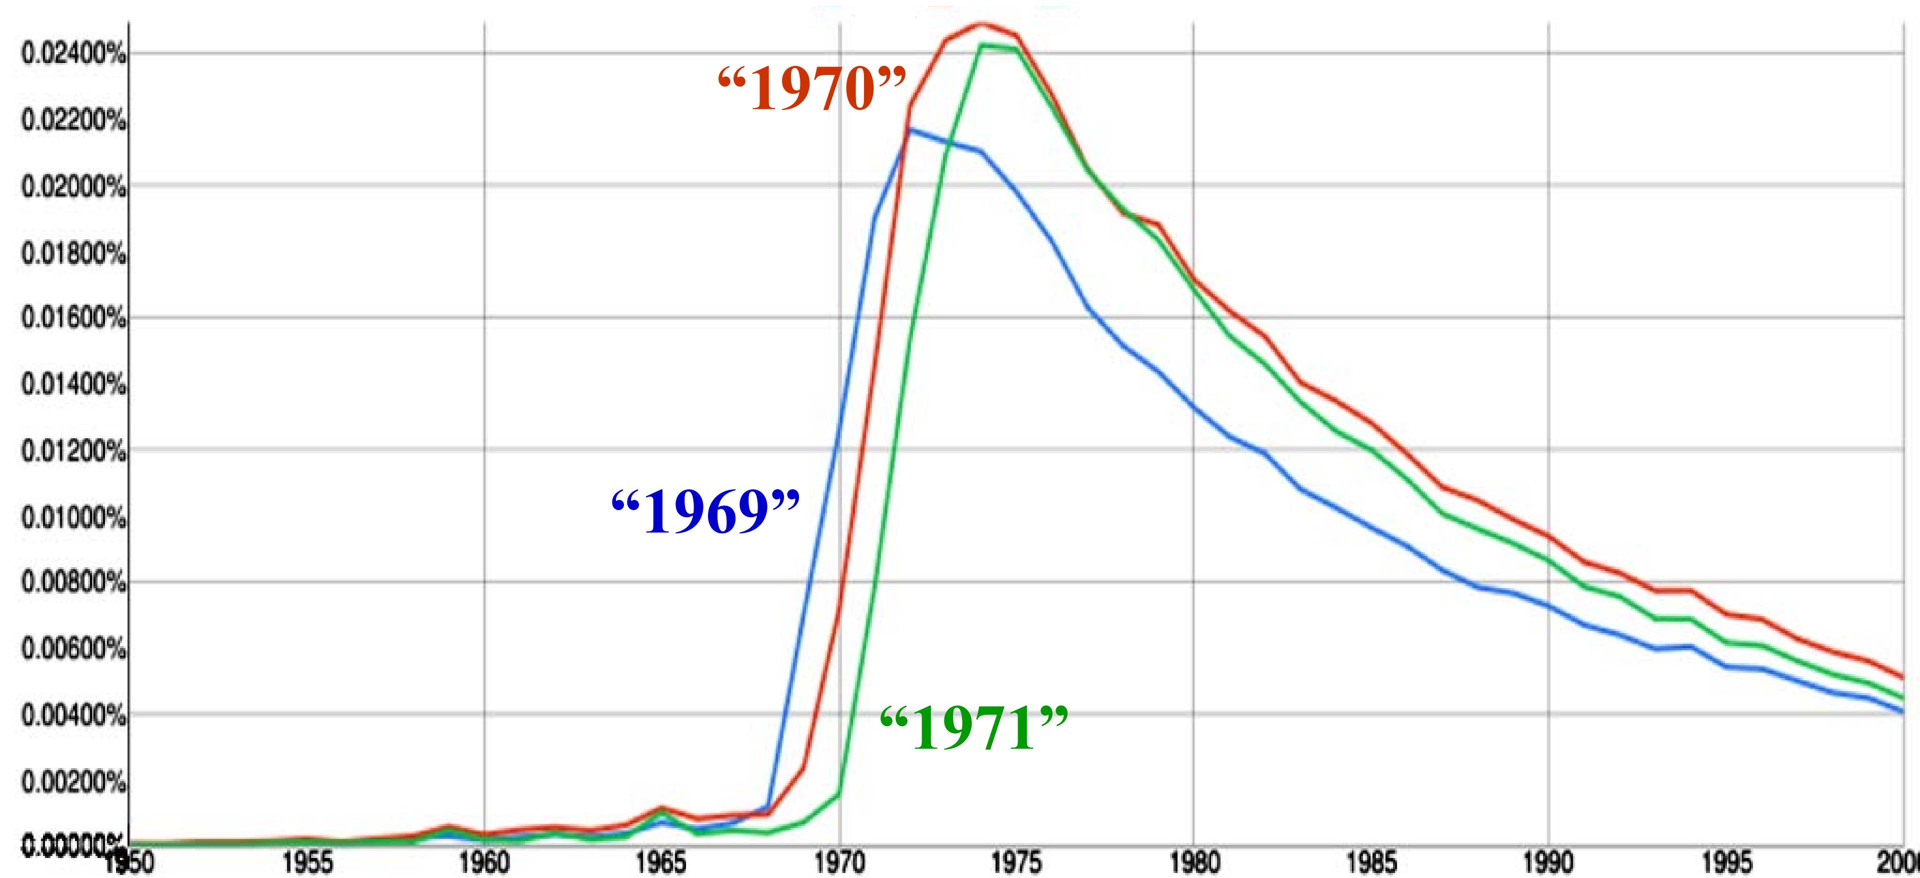

Unlike the future, past macroscopic events are generally regarded as definite and unambiguous. Of course some *microscopic* “events” in the past (e.g. which path an unobserved photon followed through an interferometer) are regarded as being ambiguous, not because of ignorance, but because they are ill-defined in principle.

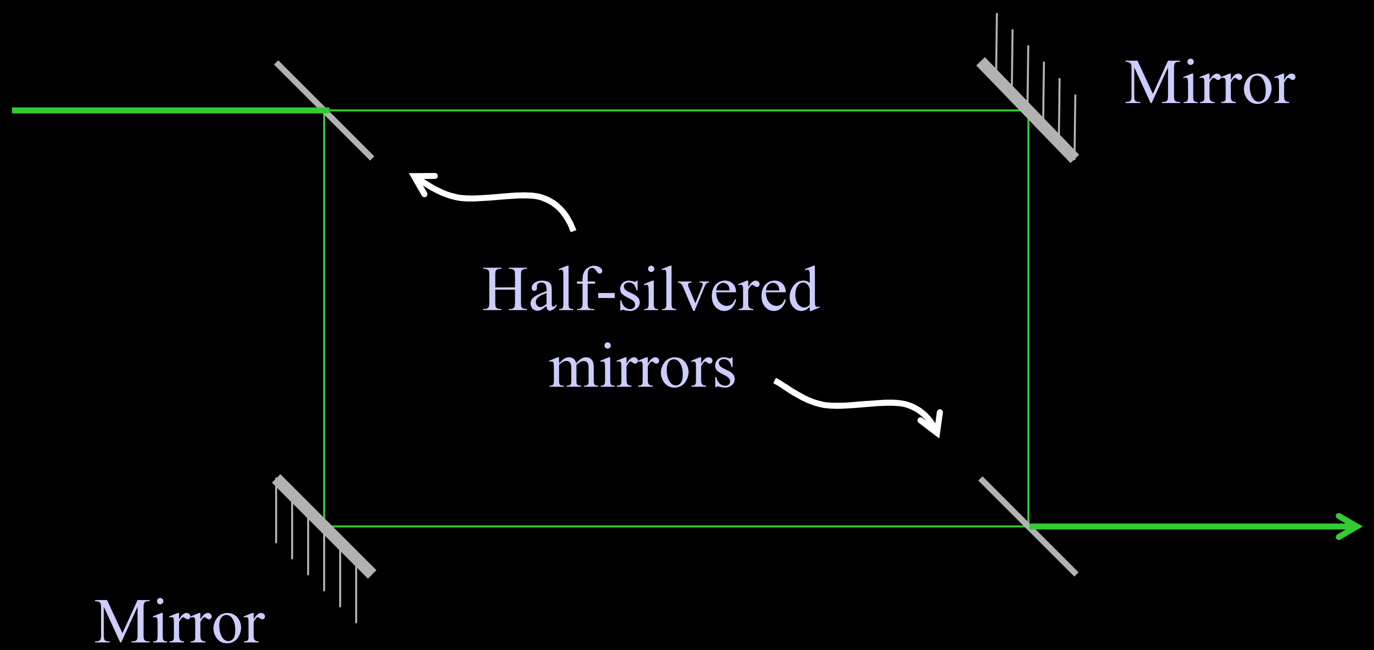

If either path through the interferometer is blocked, the photon leaves both exits equally often.

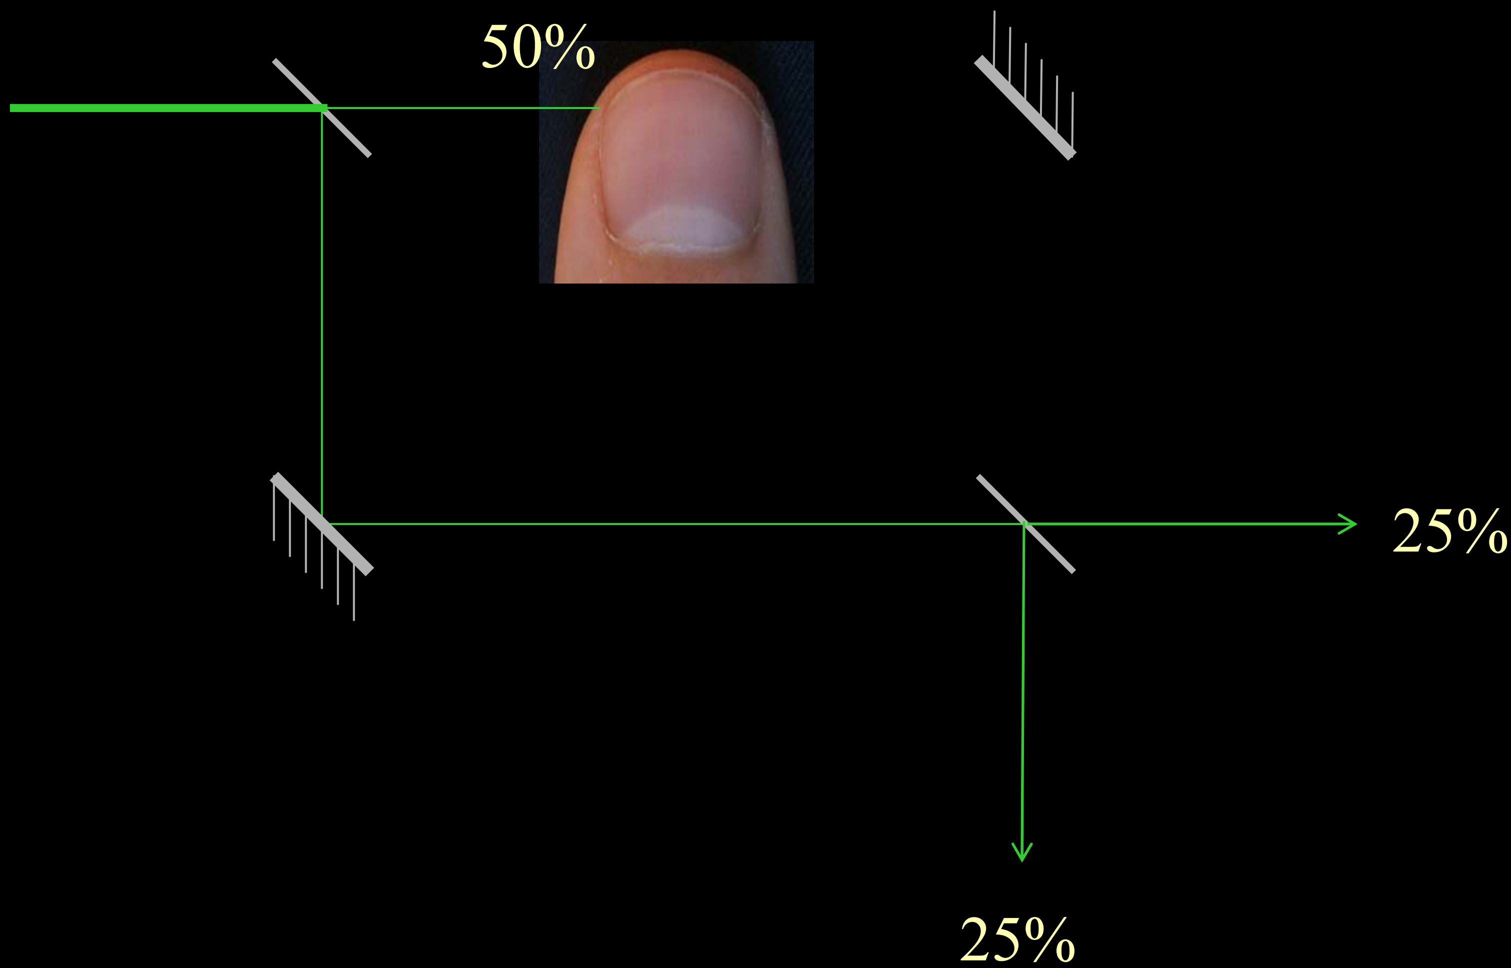

But with both paths left open, the photon always leaves by the same exit, indicating that while passing unobserved through the apparatus, it followed a **superposition** of both paths.

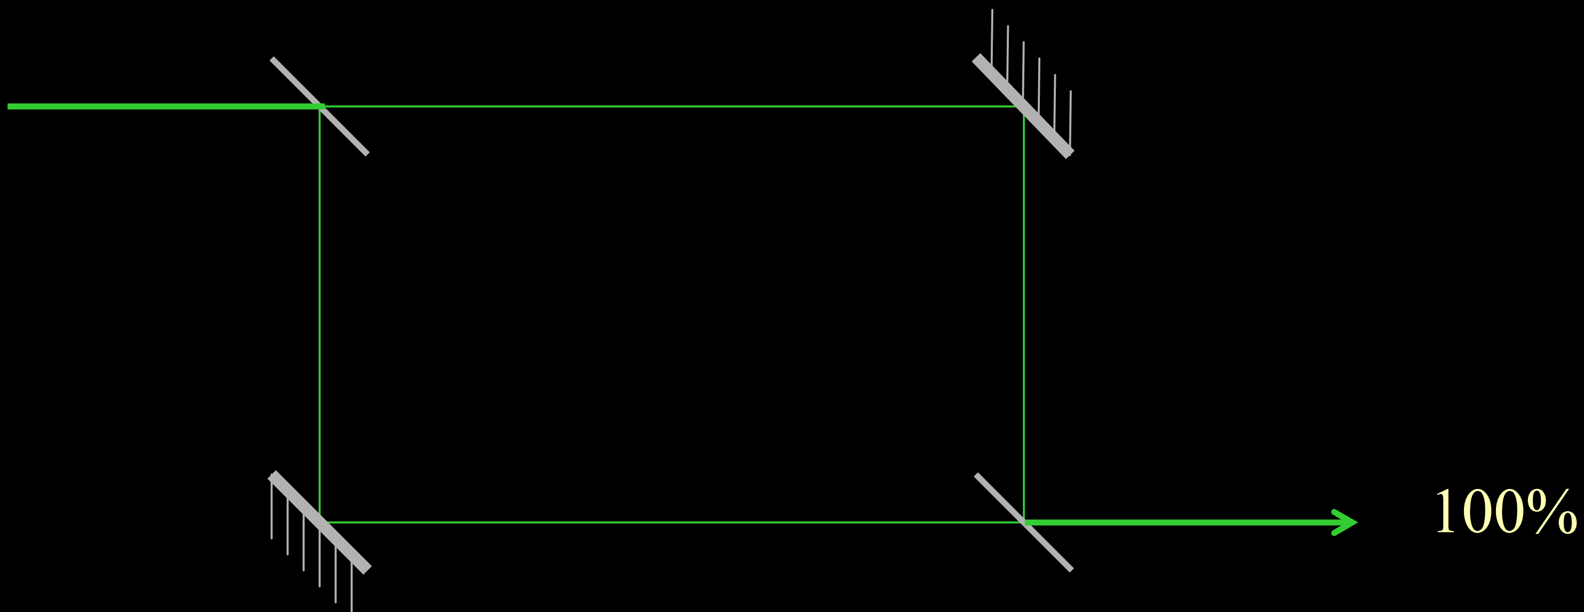

After the experiment is over, even God doesn't remember which path the photon followed.

Nowadays, it is tempting to believe that once information has become public, and someone has written about it in their blog or put it on You-tube, it can never be destroyed.

The modern world appears very different in this regard from the ancient pre-Gutenberg era, when major literary works were written down, performed, and widely known, but then then lost.

Sappho, ca 620-525 BC, by Gustav Klimmt

But even in today's world, much macroscopic, publicly accessible information is lost because no person, nor any natural process, happens to record it in a durable medium.

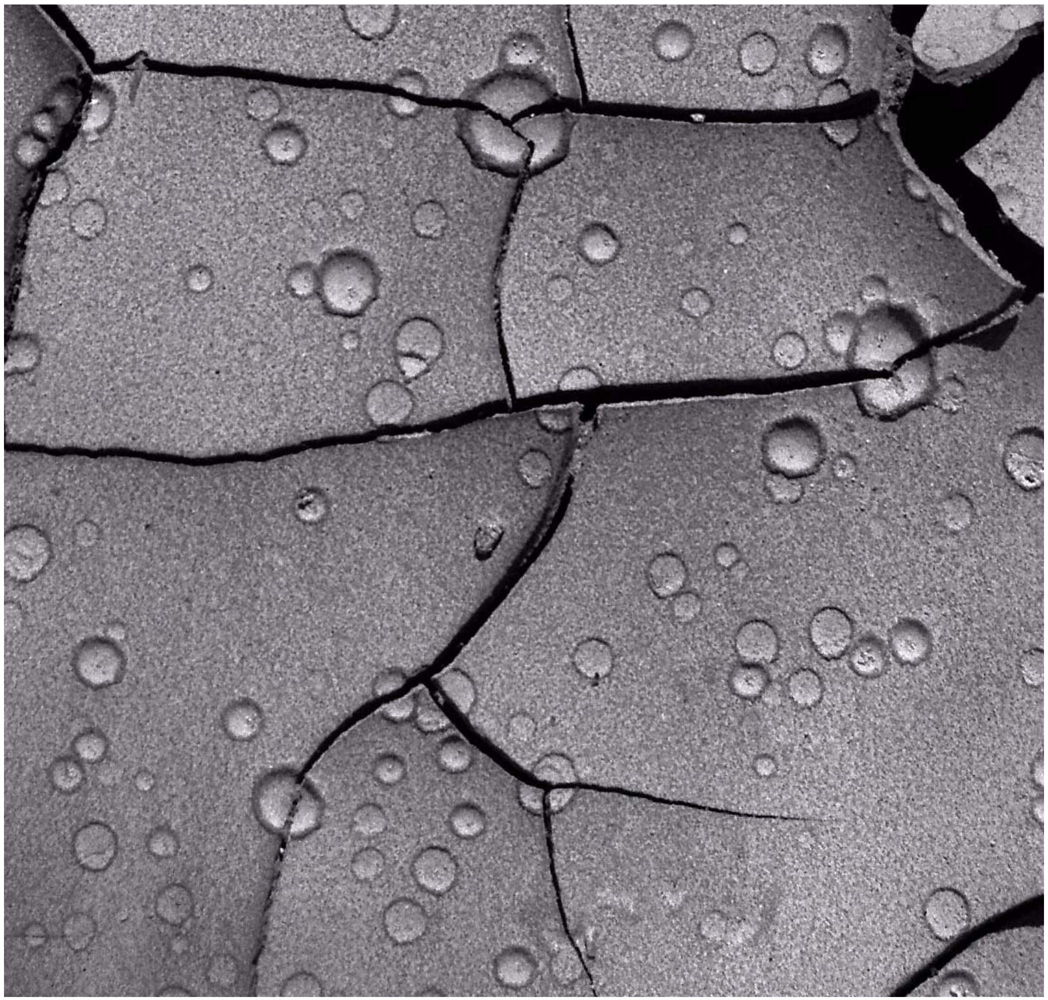

Raindrop marks in dried mud in a river bed in Las Vegas, USA in 1965.
A few days later these cracks and craters were washed away by a subsequent rain.

If no one had photographed them, would all record of them have been lost?

It is tempting to believe that such macroscopic information is not lost, just that it becomes so diffusely and complexly spread out as to be irrecoverable in practice while being stored somewhere in the universe (just as when a book is burned its contents are in principle recoverable from the exact state of the smoke, ashes, and heat it generated).

To believe otherwise is venturing dangerously close to the postmodernist view, abhorred by most scientists as arrogantly anthropocentric, that the past (or maybe even the present) has no objective reality independent of our beliefs about it, and therefore that it is pointless to inquire what "actually" happened.

But I will argue that most macroscopic, once-tangible information about the past becomes lost, not from the universe, but from the world, being swept away in the massive information flow passing through the planet.

The Earth has finite information storage capacity, but it receives a lot of entropy from the sun and exports several fold more in the form of thermal radiation into the sky.

Thermal entropy export rate $\approx$ 200 watts/sq meter at 300K $\approx 10^{30}$ bits per square meter per year.

Geological information capture rate in "hard" degrees of freedom, stable for geological times against thermal motion (e.g. atomic substitutional disorder and crystal lattice defects in solid rock of earth's crust) = crust thickness ($\approx$10 km) $\times$
rock information density ($\approx$1 bit/cubic nm) / rock lifetime ($\approx 10^8$ yr)
$\approx 10^{22}$ bits / per square meter per year.

Human digital information capture rate 100GB/person x 10^9 people who are heavy information users $\approx 10^{21}$ bits per year

(that's for the whole world, not per sq meter)

Randomizing dynamics in a representative case.

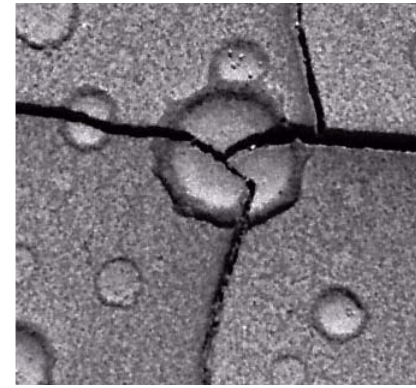

Though the raindrop originates in quantum and thermal fluctuations, it does not fall in a superposition of places. Independent observers would agree where it fell, and as long as the drop or its crater exists, reflected light will generate a torrent of replicas of the information, streaming out into space.

However, unless the crater is lucky enough to get fossilized or photographed, it will be washed away, and its former location will then lose any stable earthly embodiment. The torrent of optical replicas will cease, and the old optical replicas will escape into space. So the information of where it was remains in the universe, but not the Earth.

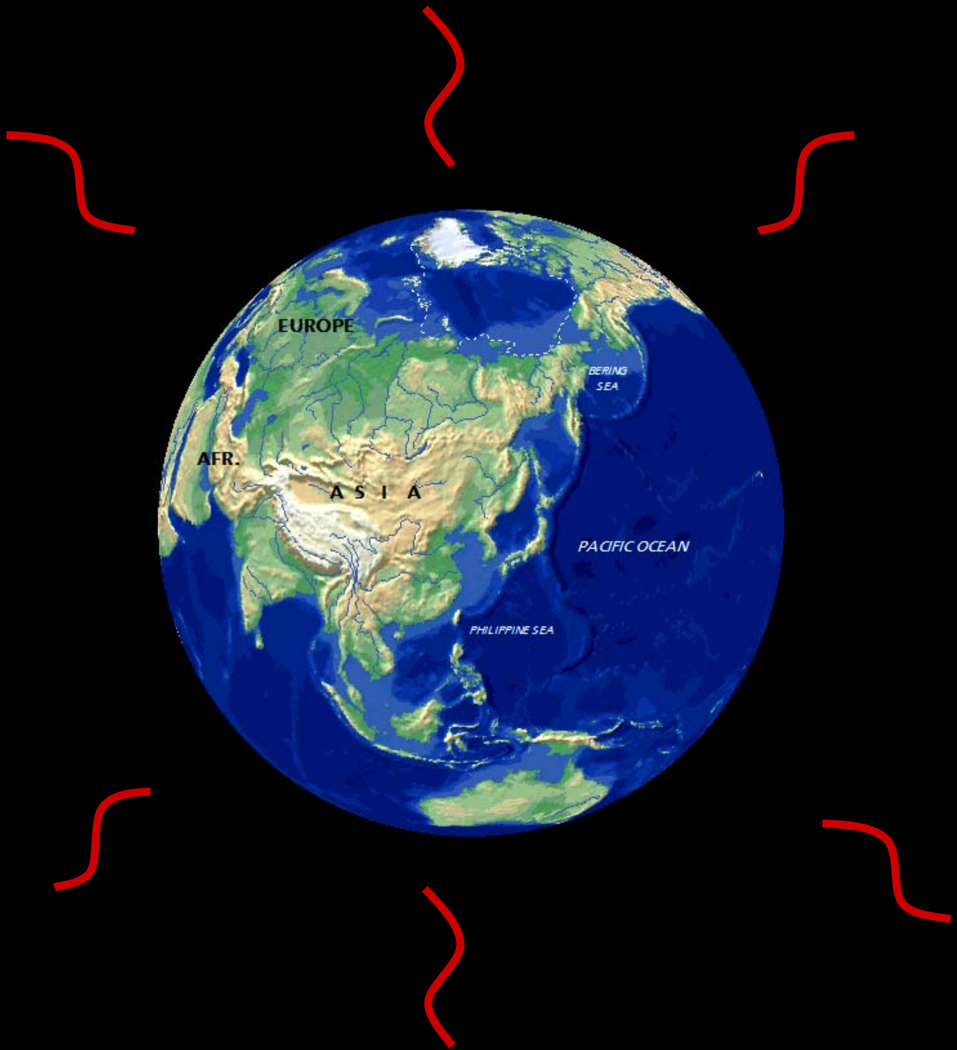

To catch up with the thermal radiation leaving Earth, one would need to travel faster than light. So the information is still in the universe, but not recoverable by us.

Mysteries of the Past:

Still recorded on earth, though unknown to any human and inaccessible with current technology:

- Locations of gold rings, dropped in an annual ceremony into the Venice Lagoon over a period of several centuries, to symbolize Venice's marriage to the Sea.

Maybe still recorded on earth, maybe escaped: Fates of mysteriously disappeared persons such as

- Physicist Ettore Majorana disappeared 1938
- Labor leader Jimmy Hoffa disappeared 1975
- Computer Scientist Jim Gray disappeared 2007

Most likely escaped:

- Unrecorded raindrops from past rain storms.
- Pattern of foam on my yesterday morning's cappuccino.

What can we do to make a particular chosen body of information long-lasting (say until the sun turns into a red giant)?

Why would we want to?
- To preserve important works of literature
- To preserve evidence of a crime until it is safe to publicize, thereby discouraging crime even in times of despotism and corruption
- Because we hate postmodernism and want to make even unimportant details of the past uncontestable.

Record the information in a durable digital medium, and bury many copies in geologically stable rock formations in various parts of the world, as if it were nuclear waste.

But suppose we wanted to store not all or most, but a lot of information, say a real-time video surveillance of entire earth surface at millimeter-millisecond resolution.

This works out to about 10^{16} bits/sq m year, well within geological capture rate.

Is this scary thing perhaps happening already, automatically, without deliberate human effort, just because frozen accidents in newly formed rock in a sense provide a **hash** of the current state of the earth?

Probably not, due to randomizing effect of dynamics. A minority share of the output of a random transformation carries negligible information about a minority share of its input.

How to obliterate earthly evidence of Jimmy Hoffa’s demise? (Former US labor leader disappeared in 1975, presumed murdered by the New York City Mafia, but body was never found. Police are still searching.)

- Cremate his body and let the smoke and heat escape

- Dissolve the ashes to make a clear liquid, with no solid fragments, then pour the liquid into the ocean

- Don’t tell anyone you did it, even on your deathbed

- For good measure, have yourself cremated and your ashes dissolved to make sure physical traces of your memory are thoroughly gone.

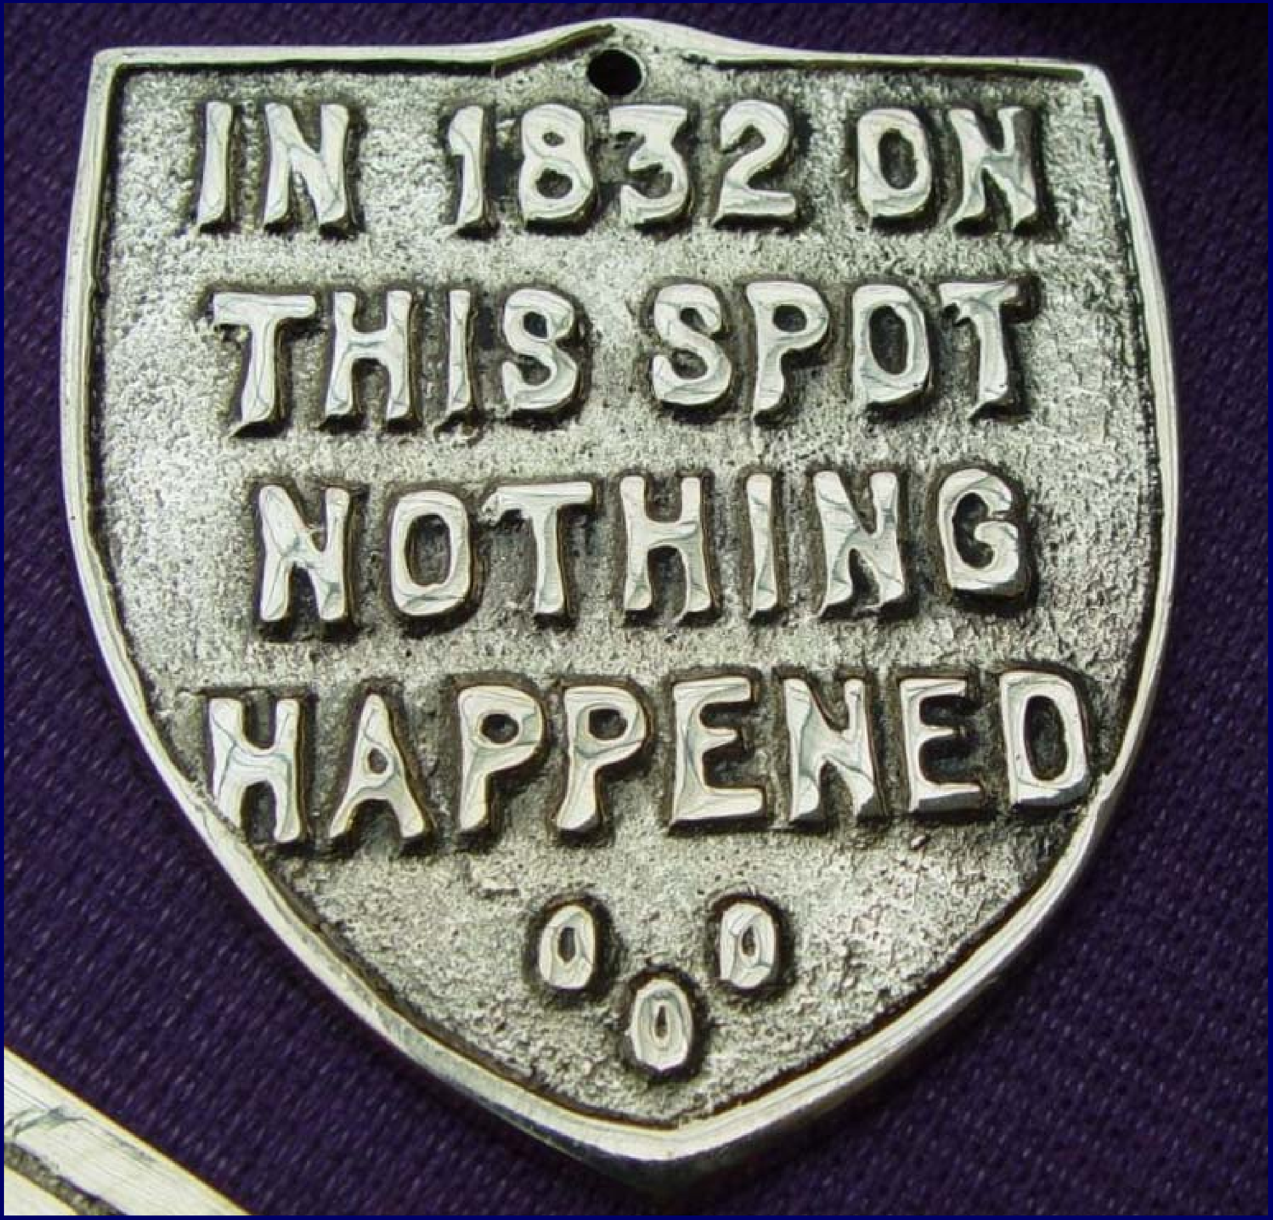
IN 1832 ON
THIS SPOT
NOTHING
HAPPENED

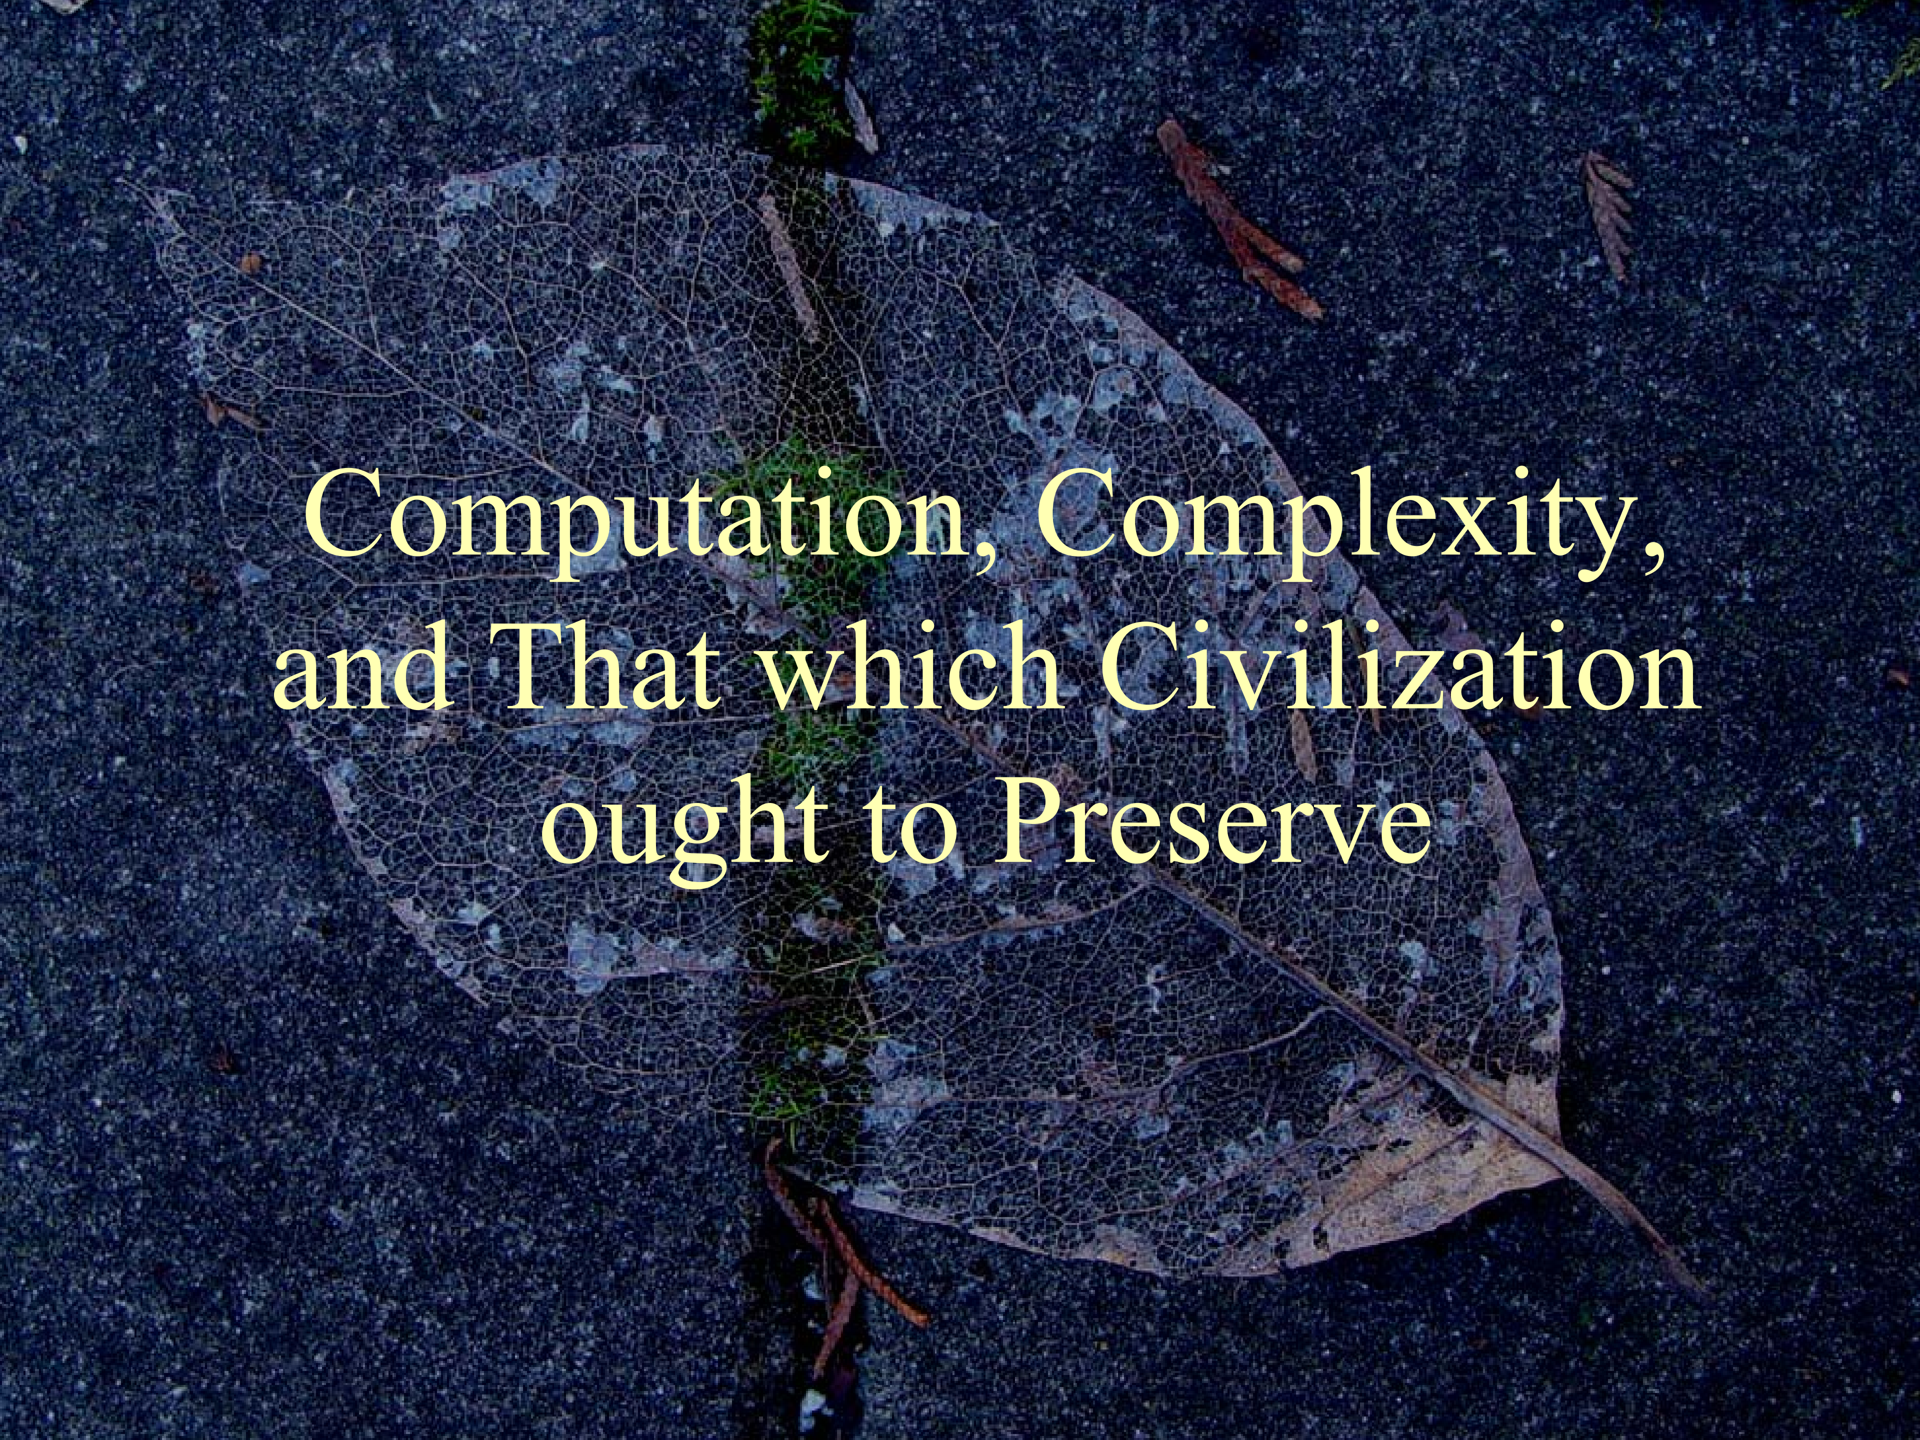
Computation, Complexity, and That which Civilization ought to Preserve

It is often said that science and ethics are separate, or that science should be subsidiary to ethics
(e.g. by having scientists avoid research with bad applications or consequences), but I want to ask:

- What positive contributions can science and math make to setting ethical values? Can they be used to arrive at less anthropocentric notions of good and evil?

- Can they help negotiate difficult policy questions, where attaining one good seems to require sacrificing another, and where the intensity of people's beliefs exceeds their ability to persuade one another?

- Security vs. Privacy — *CCTV cameras, Gov't monitoring of phone, email*
- Traditional vs. Rational/Secular —*10 Commandments vs. John Stuart Mill*
- Universal Human Rights vs. Respect for Indigenous Customs
- Natural vs. Artificial — *organic vs. genetically modified*
- Global vs. Local — *made in China vs. made in Boston*

In the US, these questions often get shoe-horned into a single left-right spectrum, an approach which Ambrose Bierce made fun of when he defined "*conservative*" as *"a statesman who is enamored of existing evils, as distinguished from the liberal, who wishes to replace them with others."*

Worldwide, the situation is more complicated: e.g. Russians typically have less respect for traditional/religious values than Americans, but are also more willing sacrifice liberty for security.

Secular / Rational

World Values Survey

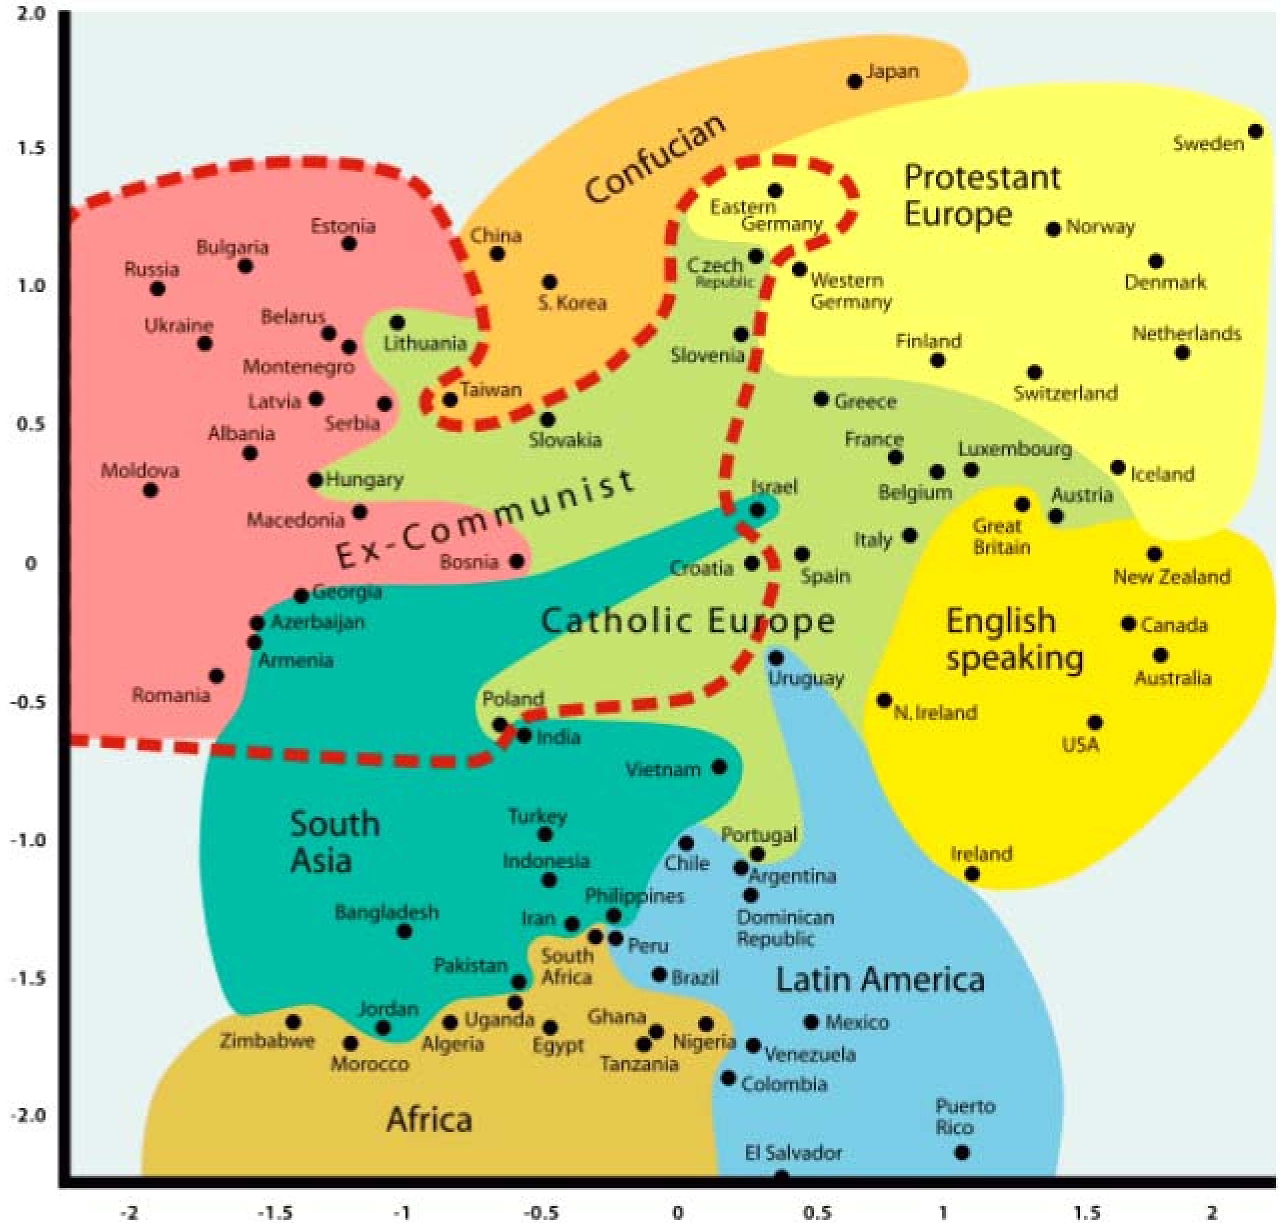

Religious / Traditional

Survival

Self-expression

Broadly, wealth is associated with liberty and secularism, but the US and Ireland are religious though rich, while post-communist societies are secularist though poor.

Should ethics be solely human-oriented?

For traditionalists, anthropocentrism, or maybe even ethnocentrism, is not a bad thing.

For rationalists any type of centrism is suspect.

Who should have rights: Foreigners? Enemy Combatants? Children? Embryos? Animals? Plants? Robots?

“Mineralarians” would avoid eating plants or animals, by subsisting on a diet of synthetic chemicals (glycerol, amino and fatty acids, vitamins, and minerals) all made from rocks, air, and water.

Can science and mathematics help provide a less anthropocentric definition of good?

That would further the secular/rational quest of purging ethics of all kinds of centrism.

One idea: use science to explore and characterize nature, and identify natural as good, artificial as bad. But we humans are certainly a product of nature, and so is everything we make. Do we really want to favor nature's older products over her newer ones?

"If God had intended us to fly, He would never have given us the railroads".

Do we want to honor the creations of all species but one?

(My wife dislikes beavers because of their human-like destructiveness, but I am fascinated by them.)

Another idea: Use mathematics to define **complexity**, and identify that which is complex, or hard to replace, as valuable. A good deed would then be one that increases the complexity of the universe, building something that would be hard to replace if destroyed.

Can we find a single, objective, mathematical definition for the kind of "we know it when we see it" complexity that is so evident in living organisms and their byproducts?

We are not seeking here a definition of **life** (a different and perhaps harder task), because nonliving objects can also be complex.

“Complexity” cannot be identified with situations of maximum order, or maximum disorder, or indeed any intermediate amount of order.

It is not a thermodynamic potential, like entropy or free energy.

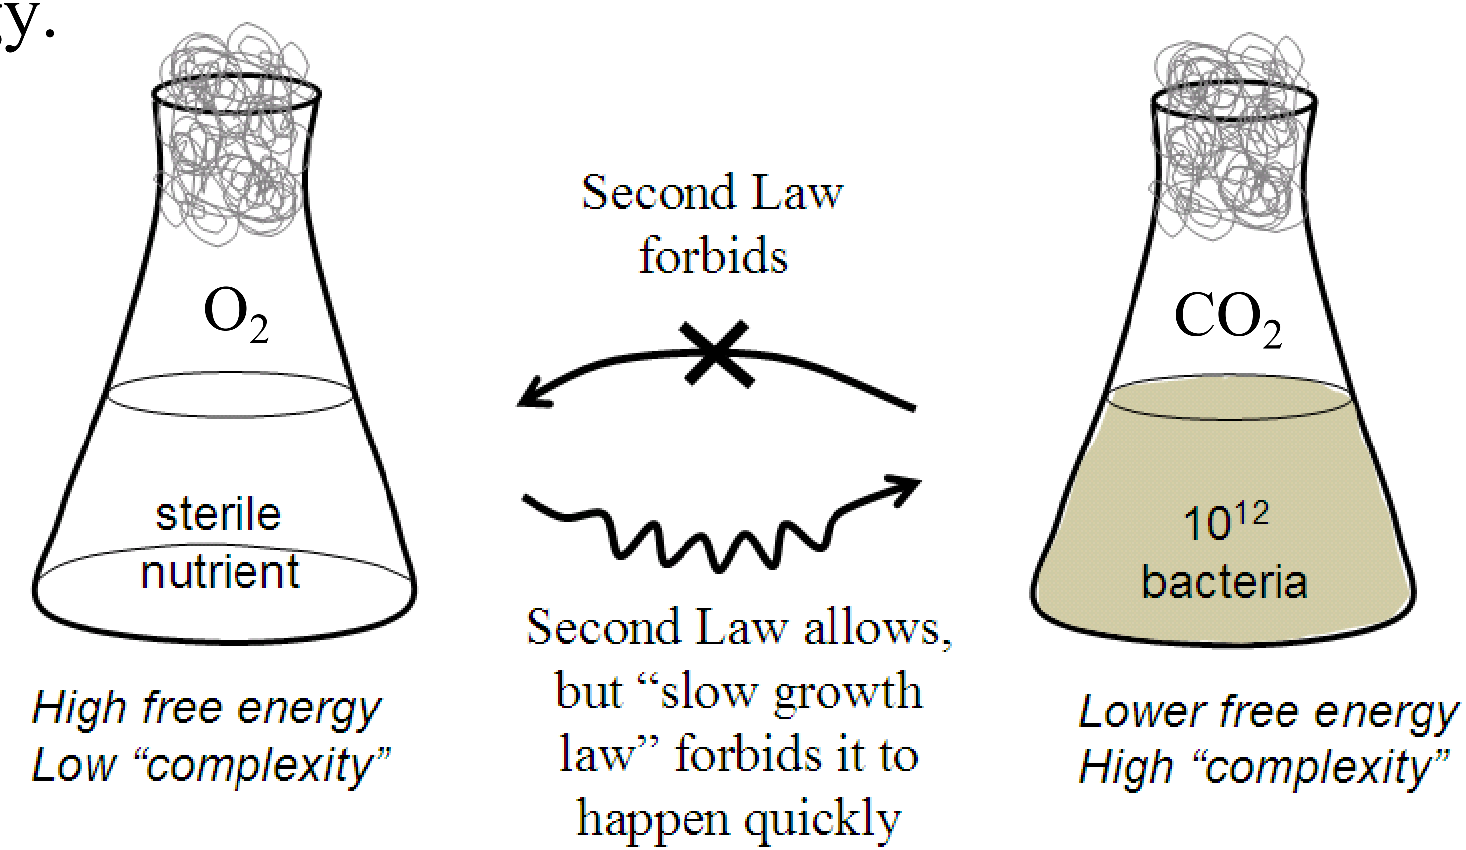

Simple dynamical processes (such as this 1 dimensional reversible cellular automaton) are easier to analyze and can produce structures of growing “complexity” from simple initial conditions.

time ⟶

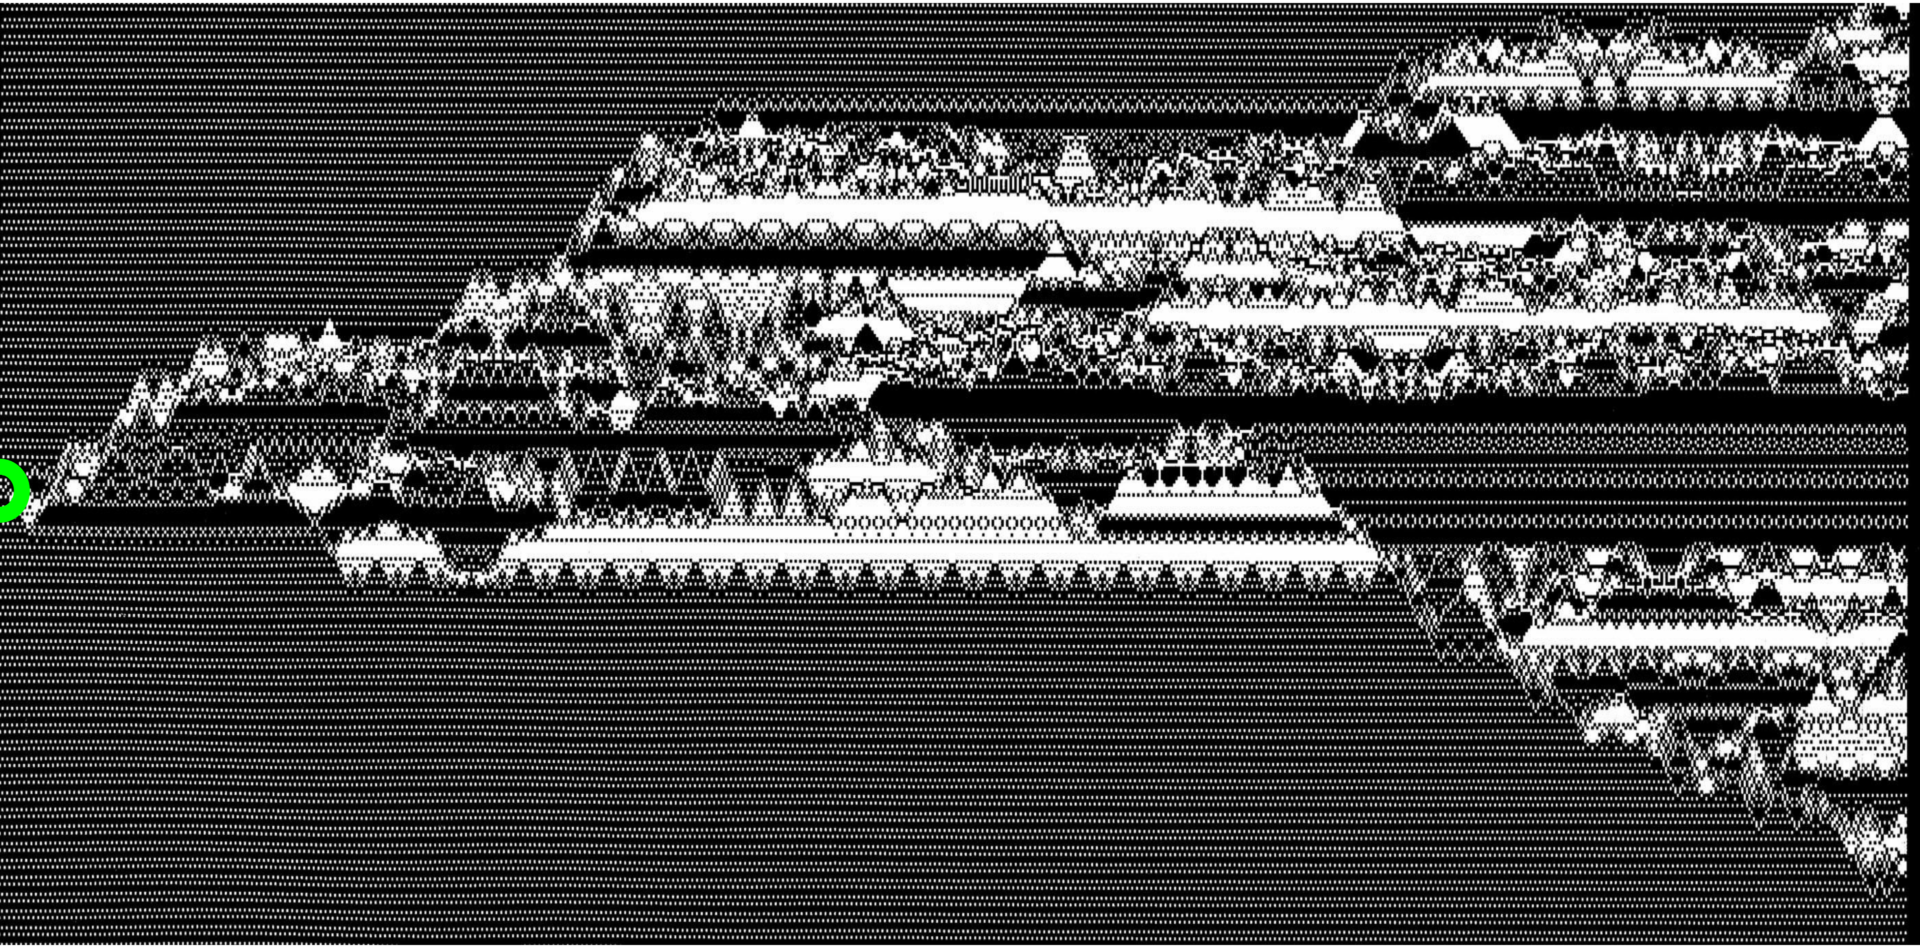

Small irregularity (green) in otherwise periodic initial condition produces a complex deterministic wake.

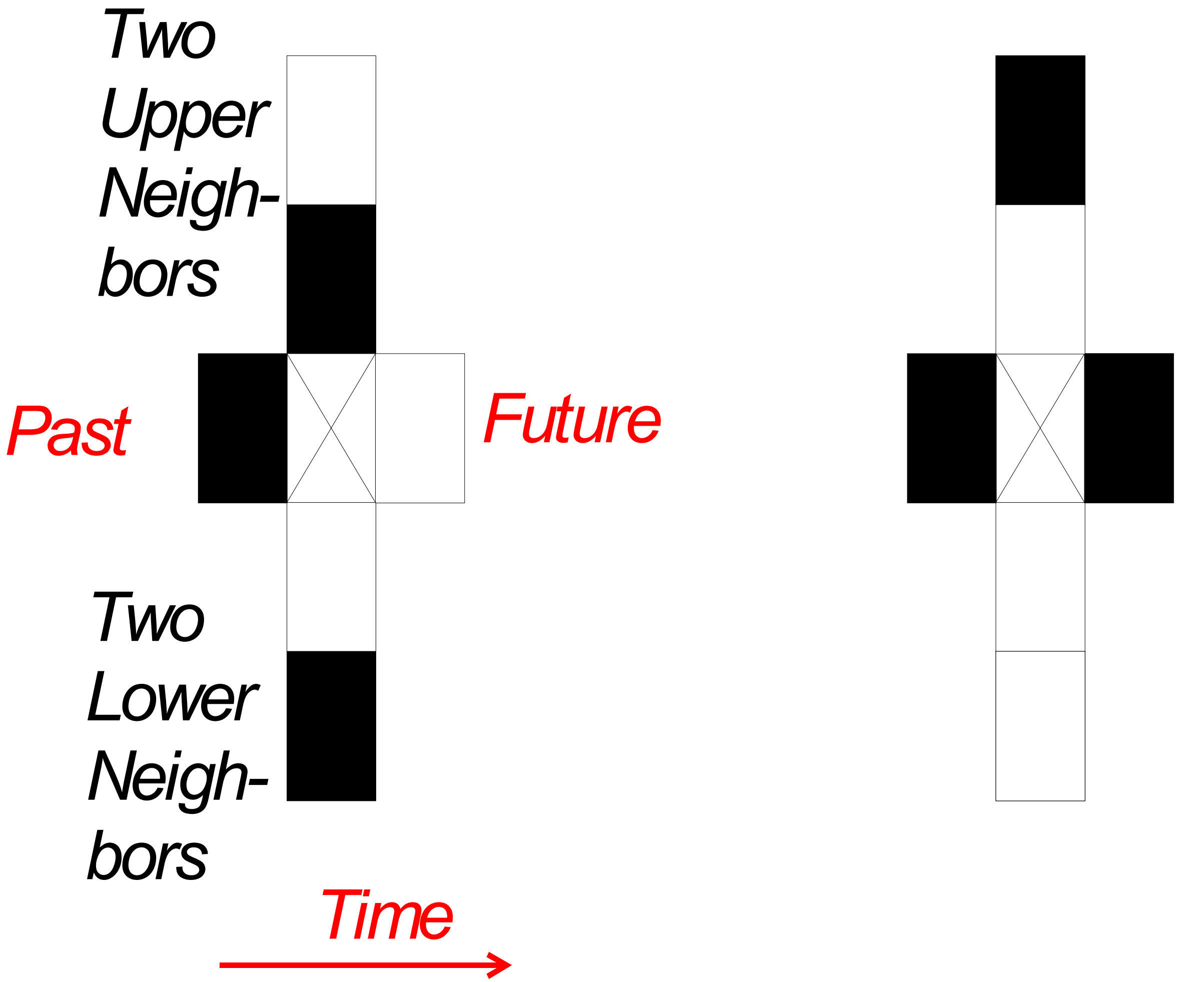

Range-2, deterministic, 1-dimensional Ising rule. Future differs from past if exactly two of the four nearest upper and lower neighbors are black and two are white at the present time.

Subjectively complicated structures typically are “logically deep,” containing evidence of a nontrivial causal history.

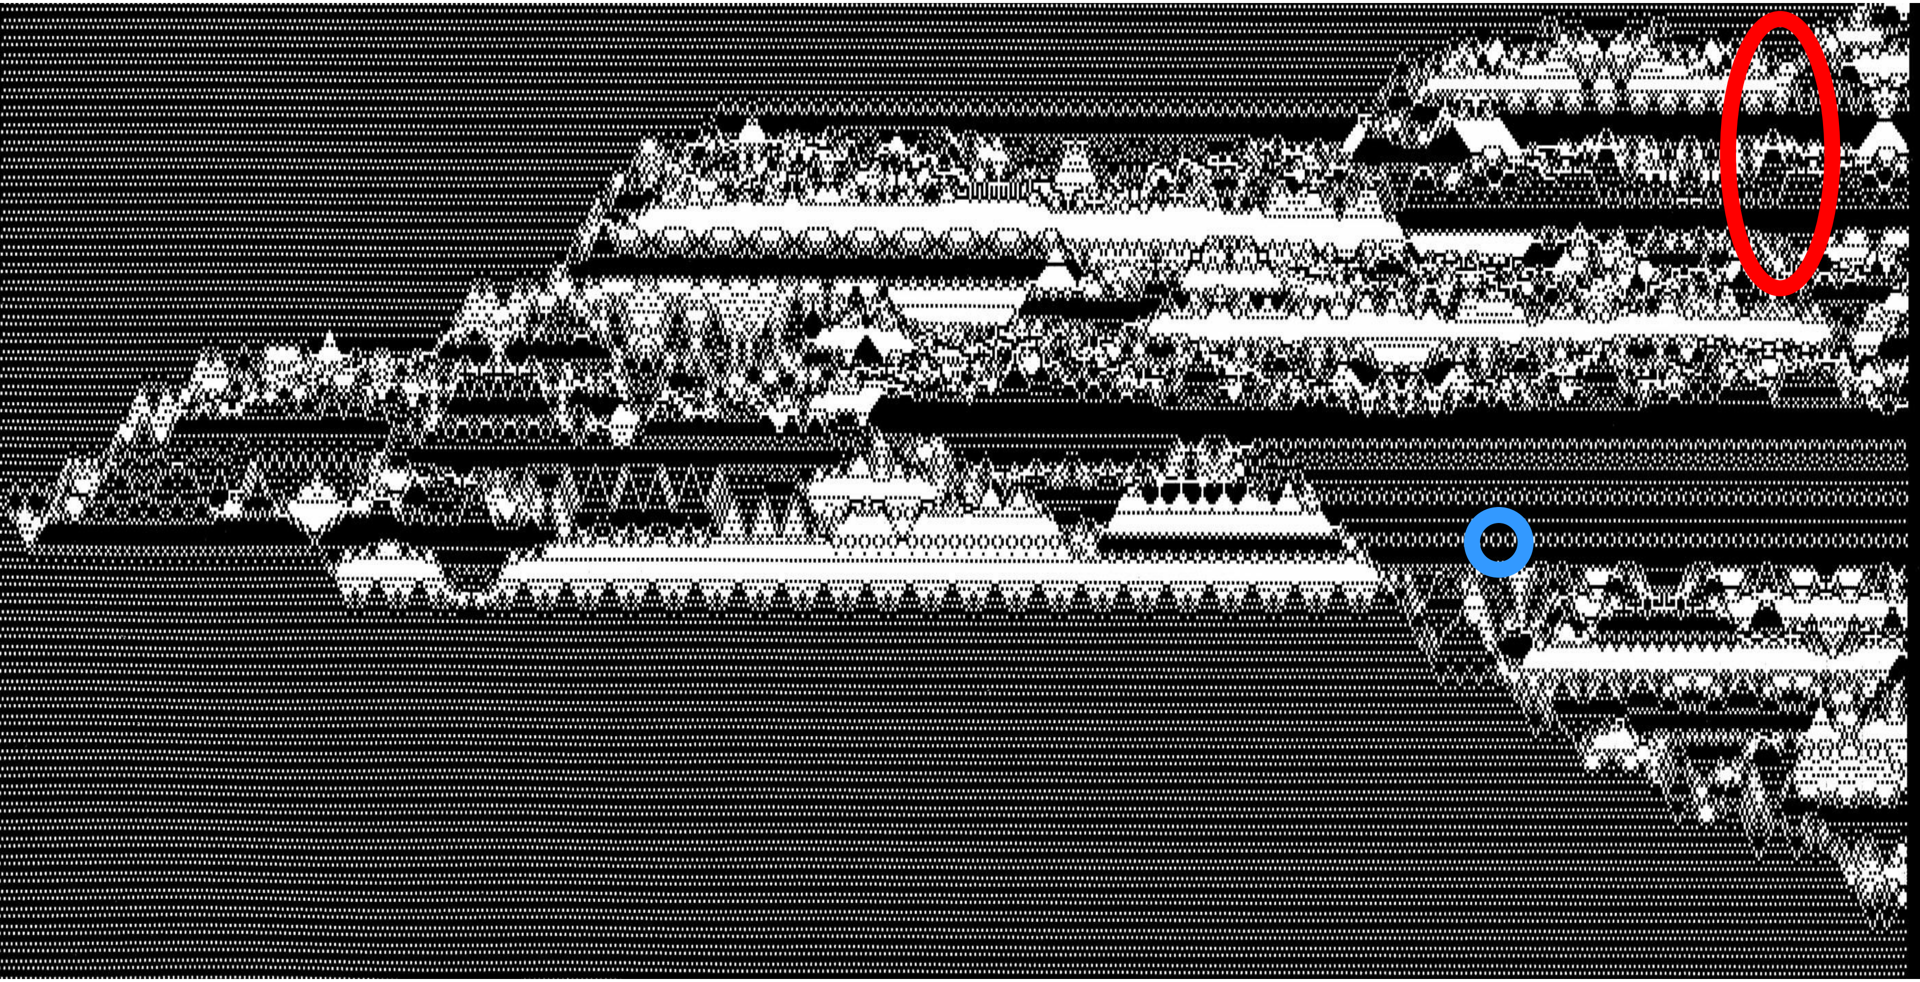

A sufficiently big piece of the wake (red) contains enough evidence to infer the whole history. A smaller pieces (blue) does not.

In the philosophy of science, the principle of Occam's Razor directs us to favor the most economical set of assumptions able to explain a given body of observational data.

Alternative hypotheses

Deductive path

Observed Phenomena

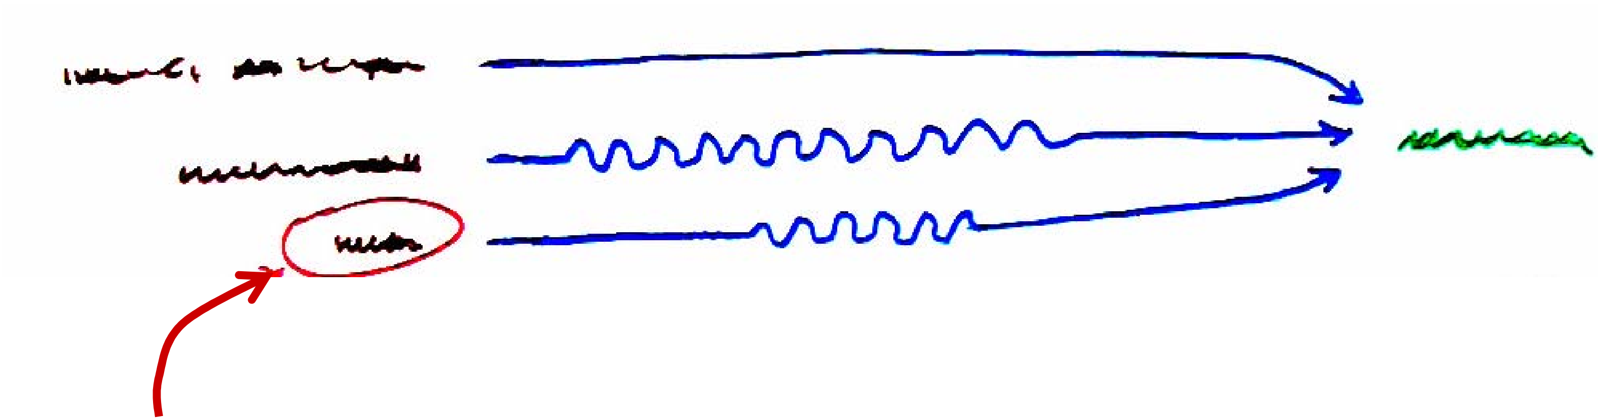

The most economical hypothesis is preferred, even if the deductive path connecting it to the phenomena it explains is long and complicated.

In a computerized version of Occam's Razor, the hypotheses are replaced by alternative programs for a universal computer to compute a particular digital or digitized object **X**.

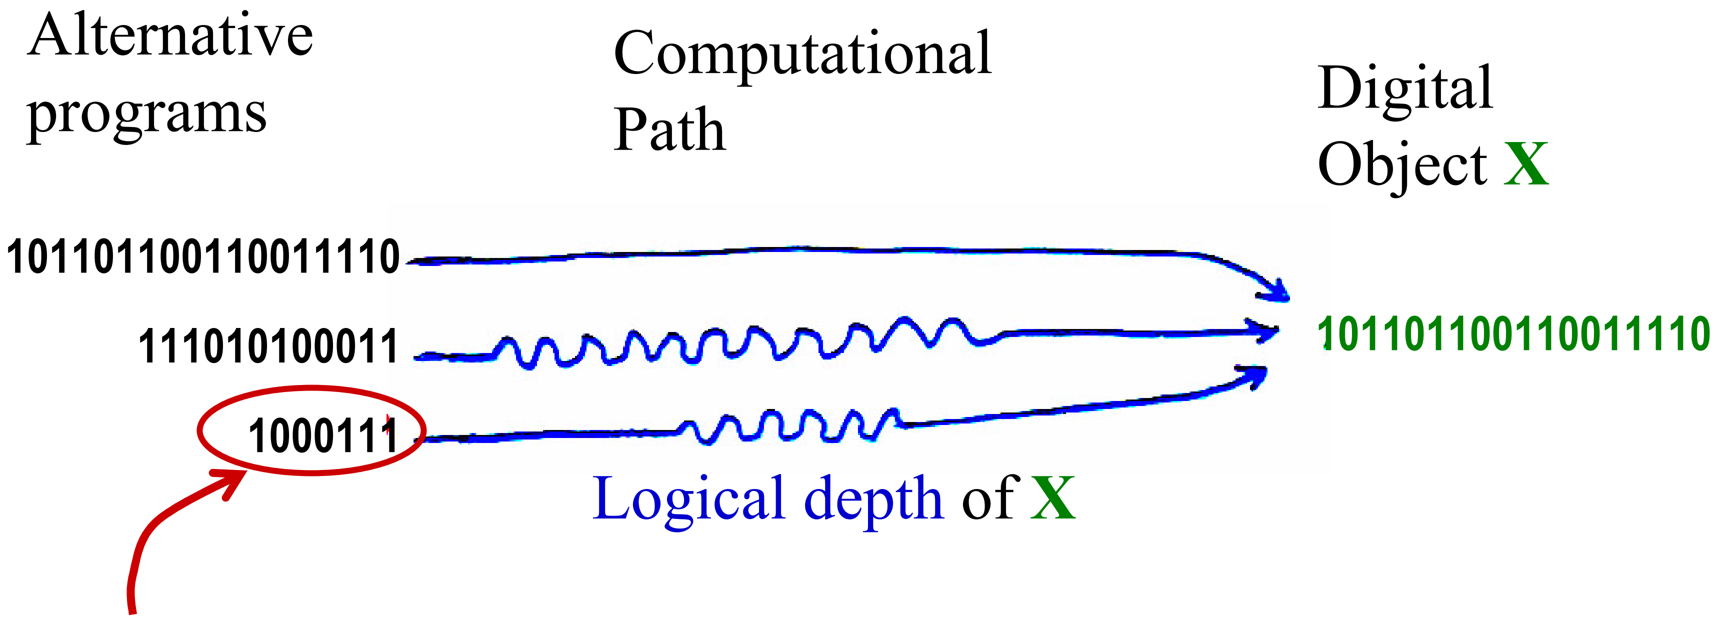

The shortest program is most plausible, so its *run time* measures the object's logical depth, or plausible amount of computational work required to create the object.

A trivially orderly sequence like 111111… is logically shallow because it can be computed rapidly from a short description.

A typical random sequence, produced by coin tossing, is also logically shallow, because it essentially **its own** shortest description, and is rapidly computable from that.

Trivial semi-orderly sequences, such as an alternating sequence of 0’s and random bits, are also shallow, since they are rapidly computable from their random part.

(Depth is thus distinct from, and can vary independently from *Kolmogorov complexity* or *algorithmic information content*, defined as the **size** of the minimal description, which is high for random sequences. A sequence’s Kolmogorov complexity measures its randomness, not its complexity in the sense intended here.)

Complexity (logical depth) as a measure of good.

Advantage: Nicely attributes value to literature, cultural artifacts, evolved genomes, ecosystems, species, and complex thoughts and emotions (treated in an utterly materialist way, as patterns of atoms in people's brains). Destroying the last copy of a good book, or the last individual of a species or human culture, is especially bad, because it destroys complex information not available elsewhere.

In many cases such losses are prevented by biological and cultural replicative processes. Even a major asteroid impact would probably not destroy all copies of Shakespeare's works, of the human genome, or all Model-T Fords.

Problem: A nuclear weapon is complex, but not good.

Possible answer: Average over the consequences. It can be argued that Hiroshima lost a lot of complexity in a few seconds, e.g. private thoughts, social relationships, and a few artifacts that were destroyed so thoroughly as to be irrecoverable, even in principle, from the rubble and the survivors' memories. This happens of course on a smaller scale when anyone dies, but in the normal course of events, a good deal of the complexity people create is salvaged in the memories of their friends and in the deeds they have done.

The good that people do lives long after them;
the evil is oft interred with their bones.

Problem: scaling with number: we would like to believe that many people's happiness is not like many copies of a good book, scarcely more valuable than one copy.

Possible answer: People's experiences are so different that they scarcely overlap. An adolescent, on falling in love for the first time, thinks "I am the first person ever to feel this." We elders smile. But for other kinds of experience it probably really is true. I find that my friends are all different, complicating my life in ways that overlap only slightly from one friend to another.

Jerusalem Talmud, Sanhedrin 4:1 (22a)

כל המאבד נפש אחת מעלין עליו כאילו איבד עולם מלא

וכל המקיים נפש אחת מעלין עליו כאילו קיים עולם מלא

Whoever destroys a life, it is considered as if he destroyed an entire world.

And whoever saves a life, it is considered as if he saved an entire world.

Question from audience: If each person is unique in their experience and complexity, why not each tree? Then when beavers kill many trees, a lot of evil has been done, even though there may be many surviving trees of the same species in the forest.

Answer: That's a good question. And if trees are unique in their experience, what about bacteria? I guess I would say that somewhere "below" humans but "above" bacteria, the individuality of experience fades away, so killing an individual is no longer like destroying a whole universe.

Another question: Some people think rocks have souls. If so, isn't it wrong to eat them?

Answer: Being made into mineralarian food and eaten is less violent than what happens to them in nature. I think the rocks would regard it as a vacation.

In Place of Laplacian Determinism, we have a world where

- Microscopic phenomena such as radioactive decay are inherently unpredictable. Some macroscopic phenomena are also unpredictable, because they depend on chaotically amplified quantum unpredictability. Next year's weather, and (at least in some countries) who will win next year's elections appear to be phenomena of this sort. Better computers and better data gathering can only make these phenomena more predictable up to limits set by their intrinsic probabilism.

- Even though the earth retains a great deal of deep information about its past, a much larger amount escapes into space, making many details of the Earth's past nearly as ambiguous as its future.

- Vastly greater computing power would lead to many magic-seeming predictions, but would not abolish unpredictability.

- Mathematically-based ethics is possible in theory, but not likely appeal to many people in practice.

Extra Slides

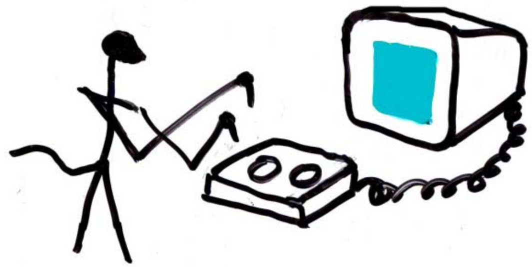

If the monkey types randomly, supplying a 0 or 1 equiprobably whenever the universal computer asks for a new input digit, the computer halts with some probability between 0 and 1. That irrational number, called Chaitin's Omega, is the total weight of all the halting branches of the Monkey Tree.

The digit sequence of Ω is algorithmically random—no computable betting strategy can make money betting against it—yet it is full of information in the sense that it can be used to generate the halting oracle ***K***.

Indeed it is a maximally concise form of this information. The first ***n*** digits of Ω suffice to decide the halting of any program of ***n*** bits or shorter, equivalent to the first 2^n bits of ***K***.

To decide the halting of an ***n***-bit program ***p***, place a weight of equal to the first ***n*** bits of Ω (a slight underestimate) on the left pan of a balance. Run all programs in parallel, and whenever one of size ***k*** bits halts, drop a weight 2^{-k} onto the right pan. When the balance eventually tips, examine all the weights in the right pan. If ***p*** is not among them, that program will never halt.

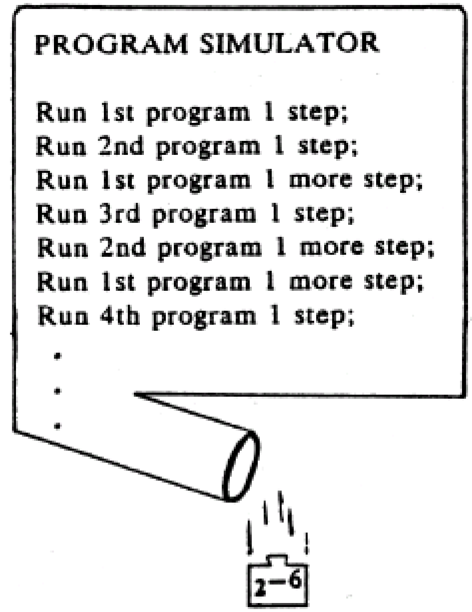

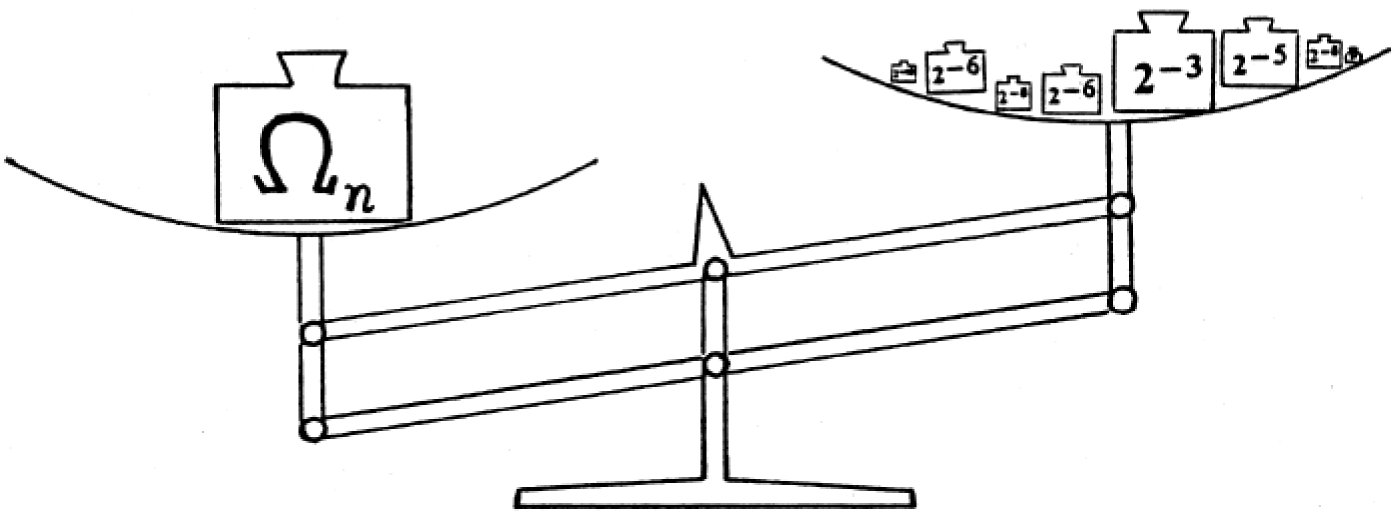

Ω is an extreme version of the irony of inscrutability. It contains all the information of ***K***, but in such a Delphic form as to be nearly useless.

To make the quantitative definition of logical depth more stable with respect small variations of the string x, and the universal machine U, a two-parameter definition is used:

A string x has depth d at significance level s if all programs p for computing x in time less than d are compressible by at least s bits, in other words, if

$$(U(p)=x \;\&\; T(p)<d) \Rightarrow (\exists q\; U(q)=p \;\&\; |q|<|p|-s).$$

Here $U(p)$ denotes the output and $T(p)$ denotes the run time of the standard universal computer U on input p, while $|p|$ denotes the length of string p in bits.

This formalizes the notion that all hypotheses for producing x in less than d steps suffer from at least s bits worth of ad-hoc assumptions.

Thus defined, depth obeys the slow-growth law.

James Gleick, *The Information, a History, a Theory, a Flood,* Pantheon Press 2011, ISBN 978-0-375-42372-7.

Gardner, Martin, "Mathematical Games" *Scientific American,* pp. 20-34, November 1979, a column based on and extensively quoting the manuscript "On Random and Hard-to-Describe Numbers" (IBM Report RC 7483) by Charles H. Bennett.
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.27.3492&rep=rep1&type=pdf

Gregory Chaitin, “A theory of program size formally identical to information theory”, *Journal of the ACM* **22**, 329-340 (1975)
http://www-2.dc.uba.ar/materias/azar/bibliografia/acm75.pdf

C.H. Bennett "Logical Depth and Physical Complexity" in
The Universal Turing Machine– a Half-Century Survey, edited
by Rolf Herken Oxford University Press 227-257, (1988)
Available at **http://bit.ly/nh0bra**

Recent blog post by C. H. Bennett on logical depth versus other complexity measures **http://dabacon.org/pontiff/?p=5912**

5-place logarithms of trigonometric functions

1°-4°
2.2418
2.2419
1.0580
.99 834
85°-88°

5°-8°
2.9403
2.9419
.80 029
.99 462
81°-84°

9°-12°
.19 433
.19 971
.83 064
.98 872
77°-80°

13°-16
.35 206
.36 330
.51 460
.98 060
73°-76

17°-20°
.46 594
.48 534
.41 582
.97 015
69°-72°

29°

′	log sin 9-10	log tan 9-10	log cot 0	log cos 9-10	′
0	.68 557	.74 375	.25 625	.94 182	60
1	580	405	595	175	59
2	603	435	565	168	58
3	625	465	535	161	57
4	648	494	506	154	56
5	.68 671	.74 524	.25 476	.94 147	55
6	694	554	446	140	54
7	716	583	417	133	53
8	739	613	387	126	52
9	762	643	357	119	51
10	.68 784	.74 673	.25 327	.94 112	50
11	807	702	298	105	49
12	829	732	268	098	48
13	852	762	238	090	47
14	875	791	209	083	46
15	.68 897	.74 821	.25 179	.94 076	45
16	920	851	149	069	44
17	942	880	120	062	43
18	965	910	090	055	42
19	987	939	061	048	41
20	.69 010	.74 969	.25 031	.94 041	40
21	032	998	002	034	39
22	055	75 028	24 972	027	38
23	077	058	942	020	37
24	100	087	913	012	36
25	.69 122	.75 117	.24 883	.94 005	35
26	144	146	854	93 998	34
27	167	176	824	991	33

30°

′	log sin 9-10	log tan 9-10	log cot 0	log cos 9-10	′
0	.69 897	.76 144	.23 856	.93 753	60
1	919	173	827	746	59
2	941	202	798	738	58
3	963	231	769	731	57
4	984	261	739	724	56
5	.70 006	.76 290	.23 710	.93 717	55
6	028	319	681	709	54
7	050	348	652	702	53
8	072	377	623	695	52
9	093	406	594	687	51
10	.70 115	.76 435	.23 565	.93 680	50
11	137	464	536	673	49
12	159	493	507	665	48
13	180	522	478	658	47
14	202	551	449	650	46
15	.70 224	.76 580	.23 420	.93 643	45
16	245	609	391	636	44
17	267	639	361	628	43
18	288	668	332	621	42
19	310	697	303	614	41
20	.70 332	.76 725	.23 275	.93 606	40
21	353	754	246	599	39
22	375	783	217	591	38
23	396	812	188	584	37
24	418	841	159	577	36
25	.70 439	.76 870	.23 130	.93 569	35
26	461	899	101	562	34
27	482	928	072	554	33