

E pluribus unum: From Complexity, Universality

Terence Tao

Abstract: In this brief survey, I discuss some examples of the fascinating phenomenon of universality in complex systems, in which universal macroscopic laws of nature emerge from a variety of different microscopic dynamics. This phenomenon is widely observed empirically, but the rigorous mathematical foundation for universality is not yet satisfactory in all cases.

Nature is a mutable cloud, which is always and never the same.

–Ralph Waldo Emerson, “History” (1841)

TERENCE TAO, a Fellow of the American Academy since 2009, is Professor of Mathematics at the University of California, Los Angeles. His publications include *An Introduction to Measure Theory* (2011) and *Solving Mathematical Problems: A Personal Perspective* (2006). Three of his books, *Structure and Randomness* (2008), *Poincaré’s Legacies* (2009), and *An Epsilon of Room* (2010), draw from the mathematical research blog he has maintained since 2007.

Modern mathematics is a powerful tool to model any number of real-world situations, whether they be natural – the motion of celestial bodies, for example, or the physical and chemical properties of a material – or man-made: for example, the stock market or the voting preferences of an electorate.¹ In principle, mathematical models can be used to study even extremely complicated systems, with many interacting components. However, in practice, only very simple systems (ones that involve only two or three interacting agents) can be solved precisely. For instance, the mathematical derivation of the spectral lines of hydrogen, with its single electron orbiting the nucleus, can be given in an undergraduate physics class; but even with the most powerful computers, a mathematical derivation of the spectral lines of sodium, with eleven electrons interacting with each other and with the nucleus, is out of reach. (The *three-body problem*, which asks to predict the motion of three masses with respect to Newton’s law of gravitation, is famously known as the only problem to have ever given Newton headaches. Unlike the two-body problem, which has a

simple mathematical solution, the three-body problem is believed not to have any simple mathematical expression for its solution, and can only be solved approximately, via numerical algorithms.) The inability to perform feasible computations on a system with many interacting components is known as the *curse of dimensionality*.

Despite this curse, a remarkable phenomenon often occurs once the number of components becomes large enough: that is, the aggregate properties of the complex system can mysteriously become predictable again, governed by simple laws of nature. Even more surprising, these macroscopic laws for the overall system are often largely *independent* of their microscopic counterparts that govern the individual components of that system. One could replace the microscopic components by completely different types of objects and obtain the same governing law at the macroscopic level. When this occurs, we say that the macroscopic law is *universal*. The universality phenomenon has been observed both empirically and mathematically in many different contexts, several of which I discuss below. In some cases, the phenomenon is well understood, but in many situations, the underlying source of universality is mysterious and remains an active area of mathematical research.

The U.S. presidential election of November 4, 2008, was a massively complicated affair. More than a hundred million voters from fifty states cast their ballots, with each voter's decision being influenced in countless ways by campaign rhetoric, media coverage, rumors, personal impressions of the candidates, or political discussions with friends and colleagues. There were millions of "swing" voters who were not firmly supporting either of the two major candidates; their final decisions

would be unpredictable and perhaps even random in some cases. The same uncertainty existed at the state level: while many states were considered safe for one candidate or the other, at least a dozen were considered "in play" and could have gone either way.

In such a situation, it would seem impossible to forecast accurately the election outcome. Sure, there were electoral polls – hundreds of them – but each poll surveyed only a few hundred or a few thousand likely voters, which is only a tiny fraction of the entire population. And the polls often fluctuated wildly and disagreed with each other; not all polls were equally reliable or unbiased, and no two polling organizations used exactly the same methodology.

Nevertheless, well before election night was over, the polls had predicted the outcome of the presidential election (and most other elections taking place that night) quite accurately. Perhaps most spectacular were the predictions of statistician Nate Silver, who used a weighted analysis of all existing polls to predict correctly the outcome of the presidential election in forty-nine out of fifty states, as well as in all of the thirty-five U.S. Senate races. (The lone exception was the presidential election in Indiana, which Silver called narrowly for McCain but which eventually favored Obama by just 0.9 percent.)

The accuracy of polling can be explained by a mathematical law known as the *law of large numbers*. Thanks to this law, we know that once the size of a random poll is large enough, the probable outcomes of that poll will converge to the actual percentage of voters who would vote for a given candidate, up to a certain accuracy, known as the *margin of error*. For instance, in a random poll of a thousand voters, the margin of error is about 3 percent.

A remarkable feature of the law of large numbers is that it is *universal*. Does the

election involve a hundred thousand voters or a hundred million voters? It doesn't matter: the margin of error for the poll will remain 3 percent. Is it a state that favors McCain to Obama 55 percent to 45 percent, or Obama to McCain 60 percent to 40 percent? Is the state a homogeneous bloc of, say, affluent white urban voters, or is the state instead a mix of voters of all incomes, races, and backgrounds? Again, it doesn't matter: the margin of error for the poll will still be 3 percent. The only factor that makes a significant difference is the size of the poll; the larger the poll, the smaller the margin of error. The immense complexity of a hundred million voters trying to decide between presidential candidates collapses to just a handful of numbers.

The law of large numbers is one of the simplest and best understood of the universal laws in mathematics and nature, but it is by no means the only one. Over the decades, many such universal laws have been found to govern the behavior of wide classes of complex systems, regardless of the components of a system or how they interact with each other.

In the case of the law of large numbers, the mathematical underpinnings of the universality phenomenon are well understood and are taught routinely in undergraduate courses on probability and statistics. However, for many other universal laws, our mathematical understanding is less complete. The question of why universal laws emerge so often in complex systems is a highly active direction of research in mathematics. In most cases, we are far from a satisfactory answer to this question, but as I discuss below, we have made some encouraging progress.

After the law of large numbers, perhaps the next most fundamental example of a universal law is the *central limit theorem*. Roughly speaking, this theorem asserts

that if one takes a statistic that is a combination of many independent and randomly fluctuating components, with no one component having a decisive influence on the whole, then that statistic will be approximately distributed according to a law called the *normal distribution* (or *Gaussian distribution*) and more popularly known as the *bell curve*. The law is universal because it holds regardless of exactly how the individual components fluctuate or how many components there are (although the accuracy of the law improves when the number of components increases). It can be seen in a staggeringly diverse range of statistics, from the incidence rate of accidents; to the variation of height, weight, or other vital statistics among a species; to the financial gains or losses caused by chance; to the velocities of the component particles of a physical system. The size, width, location, and even the units of measurement of the distribution vary from statistic to statistic, but the bell curve shape can be discerned in all cases. This convergence arises not because of any "low level" or "microscopic" connection between such diverse phenomena as car crashes, human height, trading profits, or stellar velocities, but because in all these cases the "high level" or "macroscopic" structure is the same: namely, a compound statistic formed from a combination of the small influences of many independent factors. That the macroscopic behavior of a large, complex system can be almost totally independent of its microscopic structure is the essence of universality.

The universal nature of the central limit theorem is tremendously useful in many industries, allowing them to manage what would otherwise be an intractably complex and chaotic system. With this theorem, insurers can manage the risk of, say, their car insurance policies without having to know all the complicated details of how car crashes occur; astronomers can

measure the size and location of distant galaxies without having to solve the complicated equations of celestial mechanics; electrical engineers can predict the effect of noise and interference on electronic communications without having to know exactly how this noise is generated; and so forth. The central limit theorem, though, is not completely universal; there are important cases when the theorem does not apply, giving statistics with a distribution quite different from the bell curve. (I will return to this point later.)

There are distant cousins of the central limit theorem that are universal laws for slightly different types of statistics. One example, *Benford's law*, is a universal law for the first few digits of a statistic of large magnitude, such as the population of a country or the size of an account; it gives a number of counterintuitive predictions: for instance, that any given statistic occurring in nature is more than six times as likely to start with the digit 1 than with the digit 9. Among other things, this law (which can be explained by combining the central limit theorem with the mathematical theory of logarithms) has been used to detect accounting fraud, because numbers that are made up, as opposed to those that arise naturally, often do not obey Benford's law (see Figure 1).

In a similar vein, *Zipf's law* is a universal law that governs the largest statistics in a given category, such as the largest country populations in the world or the most frequent words in the English language. It asserts that the size of a statistic is usually inversely proportional to its ranking; thus, for instance, the tenth largest statistic should be about half the size of the fifth largest statistic. (The law tends not to work so well for the top two or three statistics, but becomes more accurate after that.) Unlike the central limit theorem and Benford's law, which are fairly well understood mathematically, Zipf's law is

primarily an empirical law; it is observed in practice, but mathematicians do not have a fully satisfactory and convincing explanation for how the law comes about and why it is universal.


So far, I have discussed universal laws for individual statistics: complex numerical quantities that arise as the combination of many smaller and independent factors. But universal laws have also been found for more complicated objects than mere numerical statistics. Take, for example, the laws governing the complicated shapes and structures that arise from *phase transitions* in physics and chemistry. As we learn in high school science classes, matter comes in various states, including the three classic states of solid, liquid, and gas, but also a number of exotic states such as plasmas or superfluids. Ferromagnetic materials, such as iron, also have magnetized and non-magnetized states; other materials become electrical conductors at some temperatures and insulators at others. What state a given material is in depends on a number of factors, most notably the temperature and, in some cases, the pressure. (For some materials, the level of impurities is also relevant.) For a fixed value of the pressure, most materials tend to be in one state for one range of temperatures and in another state for another range. But when the material is at or very close to the temperature dividing these two ranges, interesting phase transitions occur. The material, which is not fully in one state or the other, tends to split into beautifully fractal shapes known as clusters, each of which embodies one or the other of the two states.

There are countless materials in existence, each with a different set of key parameters (such as the boiling point at a given pressure). There are also a large number of mathematical models that physicists and chemists use to model these

Figure 1

Histogram of the First Digits of the Populations of the 237 Countries of the World in 2010

Terence Tao




The black dots indicate the Benford's law prediction. Source: Wikipedia, http://en.wikipedia.org/wiki/File:Benfords_law_illustrated_by_world%27s_countries_population.png; used here under the Creative Commons Attribution-Share Alike license.

materials and their phase transitions, in which individual atoms or molecules are assumed to be connected to some of their neighbors by a random number of bonds, assigned according to some probabilistic rule. At the microscopic level, these models can look quite different from each other. For instance, the figures below display the small-scale structure of two typical models: a site percolation model on a hexagonal lattice (Figure 2), in which each hexagon (or site) is an abstraction of an atom or molecule randomly placed in one of two states, with the clusters being the connected regions of a single color; and a bond percolation model on a square lattice (Figure 3), in which the edges of the lattice are abstractions of molecular bonds that each have some probability of being activated, with the clusters being the connected regions given by the active bonds.

If one zooms out to look at the large-scale structure of clusters while at or near the critical value of parameters (such as temperature), the differences in microscopic structure fade away, and one begins to see a number of universal laws emerging. While the clusters have a random size and shape, they almost always have a fractal structure; thus, if one zooms in on any portion of the cluster, the resulting image more or less resembles the cluster as a whole. Basic statistics such as the number of clusters, the average size of the clusters, or the frequency with which a cluster connects two given regions of space appear to obey some specific universal laws, known as *power laws* (which are somewhat similar, though not quite the same, as Zipf's law). These laws arise in almost every mathematical model that has been put forward to explain (continuous) phase transitions

From Figure 2
Complexity, Universality
Site Percolation Model on a Hexagonal Lattice at the Critical Threshold



Source: Michael Kozdron, <http://stat.math.uregina.ca/~kozdron/Simulations/Percolation/percolation.html>; used here with permission from Michael Kozdron.


and have been observed many times in nature. As with other universal laws, the precise microscopic structure of the model or the material may affect some basic parameters, such as the phase transition temperature, but the underlying structure of the law is the same across all models and materials.

In contrast to more classical universal laws such as the central limit theorem, our understanding of the universal laws of phase transition is incomplete. Physicists have put forth some compelling heuristic arguments that explain or support many

of these laws (based on a powerful, but not fully rigorous, tool known as the *renormalization group method*), but a completely rigorous proof of these laws has not yet been obtained in all cases. This is a very active area of research; for instance, in August 2010, the Fields medal, one of the most prestigious prizes in mathematics, was awarded to Stanislav Smirnov for his breakthroughs in rigorously establishing the validity of these universal laws for some key models, such as percolation models on a triangular lattice.

Figure 3
Bond Percolation Model on a Square Lattice at the Critical Threshold

Terence Tao



Note the presence of both very small clusters and extremely large clusters. Source: Wikipedia, http://en.wikipedia.org/wiki/File:Bond_percolation_p_51.png; used here under the Creative Commons Attribution-Share Alike license.

As we near the end of our tour of universal laws, I want to consider an example of this phenomenon that is closer to my own area of research. Here, the object of study is not a single numerical statistic (as in the case of the central limit theorem) or a shape (as with phase transitions), but a discrete spectrum: a sequence of points (or numbers, or frequencies, or energy levels) spread along a line.

Perhaps the most familiar example of a discrete spectrum is the radio frequencies emitted by local radio stations; this is a sequence of frequencies in the radio portion of the electromagnetic spectrum,

which one can access by turning a radio dial. These frequencies are not evenly spaced, but usually some effort is made to keep any two station frequencies separated from each other, to reduce interference.

Another familiar example of a discrete spectrum is the spectral lines of an atomic element that come from the frequencies that the electrons in the atomic shells can absorb and emit, according to the laws of quantum mechanics. When these frequencies lie in the visible portion of the electromagnetic spectrum, they give individual elements their distinctive colors, from the blue light of argon gas (which,

confusingly, is often used in neon lamps, as pure neon emits orange-red light) to the yellow light of sodium. For simple elements, such as hydrogen, the equations of quantum mechanics can be solved relatively easily, and the spectral lines follow a regular pattern; but for heavier elements, the spectral lines become quite complicated and not easy to work out just from first principles.

An analogous, but less familiar, example of spectra comes from the scattering of neutrons off of atomic nuclei, such as the Uranium-238 nucleus. The electromagnetic and nuclear forces of a nucleus, when combined with the laws of quantum mechanics, predict that a neutron will pass through a nucleus virtually unimpeded for some energies but will bounce off that nucleus at other energies, known as scattering resonances. The internal structures of such large nuclei are so complex that it has not been possible to compute these resonances either theoretically or numerically, leaving experimental data as the only option.

These resonances have an interesting distribution; they are not independent of each other, but instead seem to obey a precise repulsion law that makes it unlikely that two adjacent resonances are too close to each other – somewhat analogous to how radio station frequencies tend to avoid being too close together, except that the former phenomenon arises from the laws of nature rather than from government regulation of the spectrum. In the 1950s, the renowned physicist and Nobel laureate Eugene Wigner investigated these resonance statistics and proposed a remarkable mathematical model to explain them, an example of what we now call a *random matrix model*. The precise mathematical details of these models are too technical to describe here, but in general, one can view such models as a large collection of masses, all connected to each other by

springs of various randomly selected strengths. Such a mechanical system will oscillate (or resonate) at a certain set of frequencies; and the Wigner hypothesis asserts that the resonances of a large atomic nucleus should resemble that of the resonances of a random matrix model. In particular, they should experience the same repulsion phenomenon. Because it is possible to rigorously prove repulsion of the frequencies of a random matrix model, a heuristic explanation can be given for the same phenomenon that is experimentally observed for nuclei.

Of course, an atomic nucleus does not actually resemble a large system of masses and springs (among other things, it is governed by the laws of quantum mechanics rather than of classical mechanics). Instead, as we have since discovered, Wigner's hypothesis is a manifestation of a universal law that governs many types of spectral lines, including those that ostensibly have little in common with atomic nuclei or random matrix models. For instance, the same spacing distribution was famously found in the waiting times between buses arriving at a bus stop in Cuernavaca, Mexico (without, as yet, a compelling explanation for why this distribution emerges in this case).

Perhaps the most unexpected demonstration of the universality of these laws came from the wholly unrelated area of *number theory*, and in particular the distribution of the prime numbers 2, 3, 5, 7, 11, and so on – the natural numbers greater than 1 that cannot be factored into smaller natural numbers. The prime numbers are distributed in an irregular fashion through the integers; but if one performs a spectral analysis of this distribution, one can discern certain long-term oscillations in the distribution (sometimes known as the music of the primes), the frequencies of which are described by a sequence of complex numbers known as

the (non-trivial) zeroes of the Riemann zeta function, first studied by Bernhard Riemann in 1859. (For this discussion, it is not important to know exactly what the Riemann zeta function is.) In principle, these numbers tell us everything we would wish to know about the primes. One of the most famous and important problems in number theory is the *Riemann hypothesis*, which asserts that these numbers all lie on a single line in the complex plane. It has many consequences in number theory and, in particular, gives many important consequences about the prime numbers. However, even the powerful Riemann hypothesis does not settle everything on this subject, in part because it does not directly say much about how the zeroes are distributed on this line. But there is extremely strong numerical evidence that these zeroes obey the same precise law that is observed in neutron scattering and in other systems; in particular, the zeroes seem to “repel” each other in a manner that matches the predictions of random matrix theory with uncanny accuracy. The formal description of this law is known as the Gaussian Unitary Ensemble (GUE) hypothesis. (The GUE is a fundamental example of a random matrix model.) Like the Riemann hypothesis, it is currently unproven, but it has powerful consequences for the distribution of the prime numbers.

The discovery of the GUE hypothesis, connecting the music of the primes and the energy levels of nuclei, occurred at the Institute for Advanced Study in 1972, and the story is legendary in mathematical circles. It concerns a chance meeting between the mathematician Hugh Montgomery, who had been working on the distribution of zeroes of the zeta function (and more specifically, on a certain statistic relating to that distribution known as the pair correlation function), and the renowned physicist Freeman Dyson. In his


book *Stalking the Riemann Hypothesis*, Terence Tao mathematician and computer scientist Dan Rockmore describes that meeting:

As Dyson recalls it, he and Montgomery had crossed paths from time to time at the Institute nursery when picking up and dropping off their children. Nevertheless, they had not been formally introduced. In spite of Dyson’s fame, Montgomery hadn’t seen any purpose in meeting him. “What will we talk about?” is what Montgomery purportedly said when brought to tea. Nevertheless, Montgomery relented and upon being introduced, the amiable physicist asked the young number theorist about his work. Montgomery began to explain his recent results on the pair correlation, and Dyson stopped him short – “Did you get this?” he asked, writing down a particular mathematical formula. Montgomery almost fell over in surprise: Dyson had written down the sinc-infused pair correlation function... Whereas Montgomery had traveled a number theorist’s road to a “prime picture” of the pair correlation, Dyson had arrived at this formula through the study of these energy levels in the mathematics of matrices.²

The chance discovery by Montgomery and Dyson – that the same universal law that governs random matrices and atomic spectra also applies to the zeta function – was given substantial numerical support by the computational work of Andrew Odlyzko beginning in the 1980s (see Figure 4). But this discovery does not mean that the primes are somehow nuclear-powered or that atomic physics is somehow driven by the prime numbers; instead, it is evidence that a single law for spectra is so universal that it is the natural end product of any number of different processes, whether from nuclear physics, random matrix models, or number theory.

The precise mechanism underlying this law has not yet been fully unearthed;

From Figure 4
 Complexity, Spacing Distribution for a Billion Zeroes of the Riemann Zeta Function, with the Corresponding
 Universality Prediction from Random Matrix Theory



Source: Andrew M. Odlyzko, "The 10^{22} -nd Zero of the Riemann Zeta Function," in *Dynamical, Spectral, and Arithmetic Zeta Functions*, Contemporary Mathematics Series, no. 290, ed. Machiel van Frankenhuysen and Michel L. Lapidus (American Mathematical Society, 2001), 139–144, <http://www.dtc.umn.edu/~odlyzko/doc/zeta.10to22.pdf>; used here with permission from Andrew Odlyzko.

in particular, we still do not have a compelling explanation, let alone a rigorous proof, of why the zeroes of the zeta function are subject to the GUE hypothesis. However, there is now a substantial body of rigorous work (including some of my own work, and including some substantial breakthroughs in just the last few years) that gives support to the universality of this hypothesis, by showing that a wide variety of random matrix models (not just the most famous model of the GUE) are all governed by essentially the same law for their spacings. At present, these demonstrations of universality have not extended to the number theoretic or physical settings, but they do give indirect support to the law being applicable in those cases.

The arguments used in this recent work are too technical to give here, but I will mention one of the key ideas, which my colleague Van Vu and I borrowed from an old proof of the central limit theorem by Jarl Lindeberg from 1922. In terms of the mechanical analogy of a system of masses and springs (mentioned above), the key strategy was to replace just one of the springs by another, randomly selected spring and to show that the distribution of the frequencies of this system did not change significantly when doing so. Applying this replacement operation to each spring in turn, one can eventually replace a given random matrix model with a completely different model while keeping the distribution mostly unchanged – which can be used to show that large class-

es of random matrix models have essentially the same distribution.

This is a very active area of research; for instance, simultaneously with Van Vu's and my work from last year, László Erdős, Benjamin Schlein, and Horng-Tzer Yau also gave a number of other demonstrations of universality for random matrix models, based on ideas from mathematical physics. The field is moving quickly, and in a few years we may have many more insights into the nature of this mysterious universal law.

There are many other universal laws of mathematics and nature; the examples I have given are only a small fraction of those that have been discovered over the years, from such diverse subjects as dynamical systems and quantum field theory. For instance, many of the macroscopic laws of physics, such as the laws of thermodynamics or the equations of fluid motion, are quite universal in nature, making the microscopic structure of the material or fluid being studied almost irrelevant, other than via some key parameters such as viscosity, compressibility, or entropy.

However, the principle of universality does have definite limitations. Take, for instance, the central limit theorem, which gives a bell curve distribution to any quantity that arises from a combination of many small and independent factors. This theorem can fail when the required hypotheses are not met. The distribution of, say, the heights of all human adults (male and female) does not obey a bell curve distribution because one single factor – gender – has so large an impact on height that it is not averaged out by all the other environmental and genetic factors that influence this statistic.

Another very important way in which the central limit fails is when the individual factors that make up a quantity do not fluctuate independently of each other, but

are instead correlated, so that they tend to rise or fall in unison. In such cases, “fat tails” (also known colloquially as “black swans”) can develop, in which the quantity moves much further from its average value than the central limit theorem would predict. This phenomenon is particularly important in financial modeling, especially when dealing with complex financial instruments such as the collateralized debt obligations (CDOs) that were formed by aggregating mortgages. As long as the mortgages behaved independently of each other, the central limit theorem could be used to model the risk of these instruments; but in the recent financial crisis (a textbook example of a black swan), this independence hypothesis broke down spectacularly, leading to significant financial losses for many holders of these obligations (and for their insurers). A mathematical model is only as strong as the assumptions behind it.

A third way in which a universal law can break down is if the system does not have enough degrees of freedom for the law to take effect. For instance, cosmologists can use universal laws of fluid mechanics to describe the motion of entire galaxies, but the motion of a single satellite under the influence of just three gravitational bodies can be far more complicated (being, quite literally, rocket science).

Another instance where the universal laws of fluid mechanics break down is at the *mesoscopic* scale: that is, larger than the microscopic scale of individual molecules, but smaller than the macroscopic scales for which universality applies. An important example of a mesoscopic fluid is the blood flowing through blood vessels; the blood cells that make up this liquid are so large that they cannot be treated merely as an ensemble of microscopic molecules, but rather as mesoscopic agents with complex behavior. Other examples of materials with interesting mesoscopic behavior

Terence Tao

include colloidal fluids (such as mud), certain types of nanomaterials, and quantum dots; it is a continuing challenge to mathematically model such materials properly.

There are also many macroscopic situations in which no universal law is known to exist, particularly in cases where the system contains human agents. The stock market is a good example: despite extremely intensive efforts, no satisfactory universal laws to describe the movement of stock prices have been discovered. (The central limit theorem, for instance, does not seem to be a good model, as discussed earlier.) One reason for this shortcoming is that any regularity discovered in the market is likely to be exploited by arbitrageurs until it disappears. For similar reasons, finding universal laws for macroeconomics appears to be a moving target; according to Goodhart's law, if an observed statistical regularity in economic data is exploited for policy purposes, it tends to collapse. (Ironically, Goodhart's law itself is arguably an example of a universal law.)

Even when universal laws do exist, it still may be practically impossible to use them to make predictions. For instance, we have universal laws for the motion of fluids, such as the Navier-Stokes equations, and these are used all the time in such tasks as weather prediction. But these equations are so complex and *unstable* that even with the most powerful computers, we are still unable to accurately predict the weather more than a week or two into the future. (By unstable, I mean that even small errors in one's measurement data, or in one's numerical computations, can lead to large fluctuations in the predicted solution of the equations.)

Hence, between the vast, macroscopic systems for which universal laws hold sway and the simple systems that can be analyzed using the fundamental laws of nature, there is a substantial middle ground of systems that are too complex for fundamental analysis but too simple to be universal – plenty of room, in short, for all the complexities of life as we know it.

ENDNOTES

¹ This essay benefited from the feedback of many readers of my blog. They commented on a draft version that (together with additional figures and links) can be read at <http://terrytao.wordpress.com/2010/09/14/a-second-draft-of-a-non-technical-article-on-universality/>.

² Dan Rockmore, *Stalking the Riemann Hypothesis: The Quest to Find the Hidden Law of Prime Numbers* (New York: Pantheon Books, 2005).