Parabola Volume 53, Issue 1 (2017)

On the density of Friedman numbers

Michael Brand[1]

1 Introduction

1.1 When is a number not a number?

Nowadays, at the age of the Internet and the smart-phone, one is constantly reminded of the fact that information of all types, from video clips to encyclopedia articles, can be (and, indeed, is) encoded in bits and bytes. A string of bits, however, can be thought of as essentially being a natural number: just add a leading "1" and read it out in binary. Putting these two facts together, we see that we are surrounded by numbers, each of which represents some piece of information.

If you are reading an electronic version of this article, the article itself is stored as a number. The software that reads it is a number. The operating system that runs the software is a number. The compiler that compiled these is a number and the log files that each of these generates as they function is a number. Interestingly, the page numbers that you see in this article are, themselves, coded as part of the file storing the paper, so they, too, are stored as a number, but it is a *different* number than the page number that you see. This is a case of a number representing another number.

While these facts are clear today, in the 21st century, in 1931, when Kurt Gödel published his Incompleteness Theorem [6], this was one of several quantum leaps that the theorem introduced: math can be described by numbers.

The way to describe math by numbers is, essentially, the way this document is stored electronically. After all, this document contains a mathematical discussion, so it is, itself, math described by a number. Gödel's insight was that if math analyses numbers and math can be described by numbers, then math can analyze math. By invoking a technique known as Quining, Gödel managed to elicit an equivalence between the math doing the describing and the math being described, leading to his conclusion of incompleteness.

Six years later, when Alan Turing proved the Halting Theorem [13], he was doing much of the same, but instead of diagonalizing over math, he was diagonalizing over software: if software can manipulate numbers and software can be described by numbers, then software can manipulate software. The Halting Theorem is the result of an equivalence elicited between the software doing the manipulation and the software being manipulated.

[1]Michael Brand is Associate Professor at the Faculty of Information Technology, Monash University. His popular maths puzzle webpage Using your Head is Permitted has just celebrated its 10 year anniversary, and his first novel has just appeared, here.

These examples belong to the peaks of modern mathematics; they are the Olympus-dwellers. Far below them, there is another class of numbers that describe themselves, but which do so solely for your amusement. These are the narcissistic numbers.

Narcissistic numbers are numbers that can be computed from their own digits. There are many types of narcissistic numbers (see, e.g., [12]), and they differ by what computation operations we allow on their digits, as part of the process of calculating back the original number.

In spectacular contrast to the central place that the works of Gödel and Turing hold in modern math, narcissistic numbers have been all but driven out. Even the great mathematician G.H. Hardy, renowned for his stance that math for its own sake is math at its best, drives a hard line against the study of narcissism. He writes [7]: "These are odd facts, very suitable for puzzle columns and likely to amuse amateurs, but there is nothing in them which appeals much to a mathematician."

And yet, occasionally, the study of narcissistic numbers does yield a result that is, clearly, mathematical rather than recreational. This paper deals with one such result. It regards one class of narcissistic numbers, known as Friedman numbers.

1.2 Meet the Friedmans

Friedman numbers are numbers that can be computed from their own digits, each digit used exactly once, by use of the four basic arithmetic operations, exponentiation and digit concatenation (as long as digit concatenation is not the only operation used). Parentheses can be used at will. An example of a Friedman number is 25, which can be represented as 5^2. An example of a non-Friedman number is any power of 10, because no power of 10 can be expressed as the result of a computation using only arithmetic operations and exponentiation if the initial arguments in the computation are a smaller power of 10 and several zeros.

Friedman numbers form sequence A036057 of the Encyclopedia of Integer Sequences [4]. They were first introduced by Erich Friedman in August 2000 [5], and the first question to be asked about them was what their density is, inside the population of integers. That is, if $F(n)$ is the number of Friedman numbers in the range $[1, n]$, what is $\lim_{n\to\infty} F(n)/n$?

If we pick a number at random from a large enough range, chances are it will be a composite number. The ratio of the number of composites to the number of primes goes to infinity for large enough ranges. By contrast, odd numbers will remain 50% of the population forever. The question asked is whether Friedman numbers vanish, statistically speaking, like primes, whether they take over the entire population, like composites, or whether they reach some middle ground, as in the case of even and odd numbers, so that neither Friedmans nor non-Friedmans ultimately vanishes.

This is a question about how difficult or easy it is for a number to describe itself and how powerful the basic operations are in this respect.

The empirical evidence was not in favor of Friedman numbers. Mike Reid, Ulrich Schimke and Philippe Fondanaiche calculated the exhaustive list of all Friedman numbers up to 10,000 (there are 72 of them), and these were later supplemented by Erich

Friedman to an exhaustive list that ranges until 100,000. In total, there are 842 Friedman numbers smaller than 100,000, or 0.842%. This is not much higher than the 0.72% among the first 10,000 numbers.

The infinite case was also studied. One infinite family of non-Friedmans was described above: the powers of 10. Infinite families of Friedman numbers also began popping up. Ron Kaminsky proved the existence of infinitely many prime Friedman numbers. Brendan Owen and Mike Reid showed (independently) that any string of digits can be the prefix to a Friedman number by appending to it a fixed suffix. (For example, 12588304 forms one such suffix. Any number N followed by the digits 12588304 is a Friedman number and can be calculated from its own digits as $N \times 10^8 + 3548^2$.) Erich Friedman augmented this result by showing that any string of digits can be the suffix to a Friedman number by preceding it with a prefix that is dependent only on the length of the suffix. (For example, if XY is a two-digit number, then $2500XY = 500^2 + XY$ is a Friedman number. More generally, 25×10^k is a prefix that can precede any k-digit suffix to make a Friedman number.)

While these results, and many like them, were able to show that the density of Friedman numbers is greater than 0, only less than one percent of the integers were previously known to be Friedman numbers. This was consistent with the empirical evidence gathered for the numbers up to $100,000$.

It was, therefore, quite a surprise when it was recently found [1] that Friedman numbers have a density of 1 within the integers. That is, that even though infinite families of non-Friedmans exist, ultimately, much like in the case of composites, if one is to pick a number randomly from a large enough range it will in all likelihood be a Friedman.

In this paper, we describe how this result is reached.

Readers wishing to know more about Friedman numbers and be updated as more results are discovered about them can check out [5], which is Erich Friedman's running blog on the topic. The results quoted above are all documented in greater detail there.

2 The toolkit

Luckily, there was no need to reinvent the basic toolkit of techniques: the great works of Gödel and Turing provide the recipe on how to elicit self-description in the form desired by us, technically known as "indirect self-reference". Two fundamental tools are used in both proofs and are reused here: tuple-encoding and Quining.

2.1 Tuple encoding

One ingredient that seems essential in creating indirect self-reference is the insight that information is not monolithic: a computer file storing a video clip is, in fact, storing a multitude of frames, each of which can be restored and displayed; the pages of this article can be retrieved individually from the file that stores the paper in its entirety; and so on. Somehow, one must provide a mechanism for representing in a single number not just one, but several items, so that these can later be unambiguously decoded.

A good choice for the method chosen for encoding depends largely on the properties of the mechanism that does the encoding and the mechanism that does the decoding, and the constraints of each. In storing the pages of this paper in a file, the computer places each page's description in a separate (and, for the most part, consecutive) set of bytes. The reason for doing so is that this makes later access to the individual pages much faster, by a machine that has byte-level read/write abilities, and which is faster reading contiguously than by random access. Turing's abstract tape machine operates under very similar constraints, so his tuple-encoding, for the Halting Theorem, followed along the same lines. By contrast, Gödel's machinery worked with arithmetic operations, so a much more convenient way for him to encode the tuple (x, y, z) was as the number $2^x 3^y 5^z$. This allowed easy encoding and decoding for Gödel's machinery, but if the frames of a video clip on your computer had been stored in the same way, decoding the video in real time would have been a technological impossibility.

For our present purposes, we also need something of a more arithmetic nature, much like Gödel's method, but actually using $2^x 3^y 5^z$ would not work. The reason for this is that such an encoding requires the use of "2", "3" and "5". These are extraneous numbers, beyond what exists in the original tuple. We, on the other hand, are looking for a method that would not require having any extraneous digits as part of our formula. Ideally, it should use only x, y and z themselves.

A better encoding for us is the following. To encode $\{x_i\}_{i=1}^t$, we use the number $\text{enc}(\{x_i\}_{i=1}^t) = x_1^{x_2^{x_3^{\cdots}}}$. The fact that the number generated in this way is exceedingly large should not bother us, as arithmetic operations work just as well on large numbers as on small ones. Far more important for us is that the encoding meets our design criterion: it refrains from using any extraneous digits. It fails only on one point: decoding is ambiguous. For example, the number 64 can encode, among other possibilities, $(2, 6)$, because $64 = 2^6$, or $(4, 3)$, because $64 = 4^3$, or $(8, 2)$, because $64 = 8^2$.

Is this a major problem? It could be, but in our case there is a workaround. We don't actually need to be able to encode a tuple of *any* numbers. We can restrict ourselves to use only a subset of the integers as integers that are encodeable as part of a tuple. As long as we can ensure that this subset of the integers is sufficiently rich, the construction will hold.

Specifically, we require of the integers inside tuples to be radical free [8]. An integer x is radical free if it cannot be represented as $x = y^z$ with $z > 1$. With this restriction, any integer can now be decoded unambiguously. For example, 64 encodes $(2, 6)$ and not any of the other possibilities. This is because $4 = 2^2$ and $8 = 2^3$ are not radical free. Even the tuple length is now unambiguously decodeable as part of this encoding (though this is a fact we will not be needing for the proof).

2.2 Quining

Indirect self reference is the ability of an entity (in this case a mathematical entity) to refer to itself without use of an explicit self-reference. A powerful technique for doing so was developed by Willard Van Orman Quine [10], and is sometimes referred to as "Quining" after him. Though it was used by both Gödel and Turing for mainstream mathematics, Quining was popularized in recreational mathematics / recreational computer science by Bratley and Millo [3] who were the first to use it in order to create a computer program that outputs its own code, a type of program that has since received the name "quine".

Typically, Quining requires several parts of the encoded entity to be repetitions of the same information, encoded in different ways. A good example of this is in Quine's original use of the technique, where he utilized it to devise what is now known as Quine's paradox:

> "Yields a falsehood when appended to its own quotation"
> yields a falsehood when appended to its own quotation.

The sentence composing Quine's paradox accomplishes self-reference by being composed of two repetitions of the same phrase ("yields a falsehood when appended to its own quotation") that receive different interpretations because one is inside quotes and the other not. In a sense, the first repetition can be viewed as "data" and the second repetition, not in quotes, can be viewed as "instructions" showing how to manipulate the data. Specifically, the instructions here are to append the data to itself, enclosing the first repetition in quotation marks.

An example for a Friedman number that works like a quine is $10411041 = 1041 \times (10^4 + 1)$. Here, too, repetition is used, with one repetition ("1041") used as data to be copied, and another repetition ("$\times(10^4+1)$") giving specific instructions regarding how to manipulate the data to form the repetitions of the original number. Note, however, that, unlike other forms of quines, Friedman-number quines do not require the various repetitions to be encoded differently in order to be interpreted differently. There is no need to place one repetition "in quotes" and the other not, in order to separate the code from the data. This is because in Friedman numbers the choice of operations used to construct the number from its own digits does not need to be encoded as part of the number. It is information extraneous to the number, and it allows us to make a different choice regarding which mathematical operations to use in conjunction with each repetition.

For our proof, we require a Quining process that generates a large number of repetitions from the original data. If the data is s and the desired number of repetitions is r, then this can be done by

$$\frac{10^{L(s)r} - 1}{10^{L(s)} - 1} s \, ,$$

where $L(s)$ is the digit length of s. We will use this $L(x)$ notation throughout.

The resulting number, whose digits repeat the digits of s a total of r times, we will denote $[s]^r$. Another way to write $[s]^r$ will be

$$\underbrace{s.s.\cdots.s}_{r \text{ times}}.$$

This uses another notation that we will use throughout: the number $x.y$ will be the number whose digits are the concatenation of the digits of x with the digits of y.

3 Outline of the proof

We wish to prove the following claim:

Theorem 1. *If $F(n)$ is the number of Friedman numbers in the range $[1, n]$, then*

$$\lim_{n\to\infty} F(n)/n = 1\,.$$

The proof itself can be divided into three steps.

In the first step, we introduce a new family of Friedman numbers. For this, let us consider, again, the families of Friedman numbers already mentioned. Consider, first, the result by Owen and Reid. They mapped out numbers, s, such that for any x it is true that $x.s$ is a Friedman number. For example, $s = 12588304$ is such a number, because

$$x.12588304 = x \times 10^8 + 3548^2\,. \tag{1}$$

We call such numbers *Friedman suffixes*, because if you add them as a suffix to any number, you get a Friedman number.

Friedman, on the other hand, showed that for any number k there is a p such that for any x of digit length k it is true that $p.x$ is a Friedman number. For example, if we pick $k = 3$, then one such p is 25000. For any x value that has exactly 3 digits $25000.x$ is a Friedman number, because

$$25000.x = x + 5000^2\,.$$

We say that p is a *Friedman prefix, when placed at position k*.

We will combine these results by constructing *Friedman infixes*. A Friedman infix is a pair of numbers, (m, k), such that any number of the form $x.m.y$ is a Friedman number, if $L(y) = k$. In fact, we require of an infix to be more general still: we would like to also support the case in which y, above, is a string of k digits beginning with leading zeros, which would technically make it not a number. To write this out, we say that for (m, k) to be a prefix we require that for any x and y and any $t < k$ we have that $x.m.[0]^t.y$ is a Friedman number, if $L(y) = k - t$.

In the first step of the proof, we will show how to generate such (m, k) pairs.

The second step of the proof is to show that this family of Friedman numbers is quite rich. For an infix (m, k), let us call $l = L(m)$ the *infix length*. What we will show is that for any given infix length we can construct a large number of (m, k)-pairs with distinct k values. In fact, for each value of l we utilize only a single m, and even so we demonstrate that the number of (m, k)-pairs increases exponentially with l.

In the third step, we show that the rate at which the number of (m, k)-pairs increases with l is sufficient to prove that $F(n)/n$ converges to 1, completing the proof of the theorem.

4 Step 1: Constructing a basic (m, k)-pair

Let us call the set of nonnegative integers that can be produced from the digits of a number, x, by use of the four basic arithmetic operations, exponentiation and digit concatenation by the name $\text{span}(x)$. (It is tempting, at this point, to redefine Friedman numbers as the set of integers, x, for which $x \in \text{span}(x)$. However, such a redefinition would not be correct: $\text{span}(x)$ includes x for every x, because x can be produced from itself by concatenating all of its own digits in order. In the definition of Friedman numbers, such concatenation-only calculations are disallowed explicitly.)

In our construction, we will use m values of the form $[s]^r$. This is because we want to use Quining in order to generate self-reference.

Consider the span of such $[s]^r$ repetitions. It is not difficult to see that for any choice of $s > 0$ one can describe any number i by the digits of $[s]^r$, if one is only given a suitable r. Here's one way to do it.

If $i = 0$, use $i = s - s$. Otherwise, calculate i as

$$\underbrace{\frac{s}{s} + \cdots + \frac{s}{s}}_{i \text{ times}}.$$

This is not necessarily the smallest possible r for any particular i.

Armed with this knowledge, it is possible to construct (m, k) where m is of the form $[s]^r$ and where s can be any number greater than 2. Here's how to do this.

First, let us assemble some basic ingredients. These include the numbers 0, 1, 10 and $L(s)$, all of which are constants and therefore all of which can be computed from repetitions of s. Let's call the number of repetitions needed for this assembly c_0, c_1, c_{10} and $c_{L(s)}$, respectively. We will also use c_k to denote some number satisfying that $k \in \text{span}([s]^{c_k})$. Though we haven't chosen k, yet, for any choice of k there will certainly be possible choices for c_k.

Next, we will pick r to be a multiple of s. The reason to do this is that if $r = r's$, then $r \in \text{span}\big([s]^{r'}\big)$ holds, because r can be represented as the summation of r' copies of s:

$$r = \underbrace{s + \cdots + s}_{r' \text{ times}}.$$

Now, let's put it all together to make a Quine, using the Quining technique from Section 2.2:

$$\left(x \times (0 + 10)^{L(s)r} + \frac{10^{L(s)r} - 1}{10^{L(s)} - 1} s\right) \times 10^k + \underbrace{0 + \cdots + 0}_{t \text{ times}} + y. \tag{2}$$

(The reason we require the form $(0+10)^{L(s)r}$ in the expression is in order to enable its replacement with $0 \times 10^{L(s)r}$ for the case $x = 0$.)

This calculation certainly results in $x.[s]^r.[0]^t.y$, but does it use each one of its digits exactly once, proving it to be a Friedman? To answer this, let us count which digits we have already used, and how many of each.

Clearly, we have made exactly one use of each digit of x, of y and of t zeros. Additionally, we have used $c_0 + 2c_1 + 4c_{10} + 3c_{L(s)} + 1 + 2r' + c_k$ copies of s. What we need is for $c_0 + 2c_1 + 4c_{10} + 3c_{L(s)} + 1 + 2r' + c_k$ to equal r. This we can do by reversing the order in which we have been working so far: instead of first picking k and then calculating a suitable c_k, let us pick c_k so as to make the equation hold, then pick as k any value in $\mathrm{span}([s]^{c_k})$. For this construction to work, only one piece of the puzzle is missing: we need to show that the appropriate value for c_k is positive. That is, we want to show

$$c_0 + 2c_1 + 4c_{10} + 3c_{L(s)} + 1 + 2r' < r\,. \tag{3}$$

Let $C = c_0 + 2c_1 + 4c_{10} + 3c_{L(s)} + 1$. By definition, $r = r's$, so we can rearrange the inequality to get $r' > \frac{C}{s-2}$. Now, all we need to do is to pick a sufficiently large r', and we're done.

5 Step 2: Finding a large number of Friedman numbers with a given infix length

In the previous section, we have shown a generic method of generating Friedman numbers based on an infix. This method tended to generate quite large Friedman numbers. For example, one can use this method to show that $([3]^{90}, 3)$ is a Friedman infix, and $[3]^{90}.[0]^3$ is a Friedman number. This number, among the smallest producible by the infix method, is almost on the order of a googol, which is to say that it is a number that is still possible to write down, but far beyond any magnitude that we can comprehend meaningfully. Clearly, there are a great many Friedman numbers out there that the infix method cannot generate.

This begs the question of whether, in the process of narrowing our discussion to just this family of Friedmans, we have not already thrown too many of them away to reach our goal density of 1. The surprising truth is that we threw away only a vanishing proportion of Friedmans, and, to drive the point further, in the next step we will narrow the discussion even more and will throw away so many of them that the smallest of the remaining Friedmans will be so large it will be impossible to write it down in any conventional way.

For this, however, let us introduce some new notation. So far, we have considered the ability to generate a number from a number, but, technically, what we needed was to generate, for example, from $[s]^r$ three copies of $L(s)$, four copies of 10, etc.. It is therefore natural to also consider the ability to generate a tuple of numbers from a number. We say that $(L(s), L(s), L(s), 10, 10, \ldots) \in \mathrm{tuplespan}([s]^r)$.

Formally, let the *tuplespan* of an integer, s, be the set of tuples

$$(x_1, \ldots, x_n)$$

such that $x_i \in \mathrm{span}(s_i)$ for each $i = 1, \ldots, n$, where $\{s_i\}_{i=1}^n$ are numbers that result from a partitioning of the digits of s.

Next, we will use a tool from coding theory. This is a reasonable move, because coding theory deals with optimal ways for coding information. In our case, we wish to encode a number by its own digits.

Let $N(s)$ be the size of the largest subset of $\mathrm{tuplespan}(s)$ in which each tuple is composed only of radical-free numbers and no two tuples are prefixes of each other, in the sense that if $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_m)$ are both elements of the subset, with $n < m$, then $x_i \neq y_i$ for at least one of the values $i = 1, \ldots, n$. We say, borrowing terminology from coding theory, that the subset is a *prefix code*.

We began by narrowing the discussion from all the Friedmans to just those Friedmans that have an (m, k) infix, where $m = [s]^r$ with $s > 2$ and $r = r's$. Let us now look at just those remaining Friedmans for which $s > 3$ and k is a multiple of $L(m)$.

Claim 1. *The number of remaining Friedman numbers for any given s is in $\Omega(g^r)$, where $g = N(s)^{\frac{s-3}{s}}$. (This is just a fancy way of saying that there exists a positive constant C_0 such that, for any large enough r, $s \geq C_0 g^r$.)*

Proof. The main restriction introduced here is the fact that k must now be a multiple of $L(m)$. We can write $k = k'L(m)$. Recall also that because $m = [s]^{r's}$ we have $L(m) = L(s)r's$. To introduce this restriction, we augment Equation (2), by replacing k in it by

$$k' \times L(s) \times (\underbrace{s + \cdots + s}_{r' \text{ times}}) \,.$$

If we follow, again, the reasoning that led us to Equation (3), but with the new formula, we now get

$$c_0 + 2c_1 + 4c_{10} + 4c_{L(s)} + 1 + 3r' + c_{k'} = r \,,$$

where k' should be chosen from $\mathrm{span}(c_{k'})$. From this we know that choosing any $k' \in \mathrm{span}\big([s]^{(s-3)r'-C'}\big)$ yields a Friedman number, where $C' = c_0 + 2c_1 + 4c_{10} + 4c_{L(s)} + 1$.

To prove Claim 1, we must show that

$$\lim_{r \to \infty} \left|\mathrm{span}\Big([s]^{(s-3)r'-C'}\Big)\right| g^{-r} > 0 \,.$$

To show this, consider once again the tuple-encoding procedure $\mathrm{enc}(\cdot)$ introduced in Section 2.1. It ensures that any tuple of radical-free numbers can be encoded uniquely. Specifically, $\mathrm{tuplespan}(s)$ includes $N(s)$ distinct tuples composed of radical-free integers that form a prefix code, and consequently $\mathrm{tuplespan}\big([s]^{(s-3)r'-C'}\big)$ includes the concatenation of any $(s-3)r' - C'$ of these tuples, creating a total of $N(s)^{(s-3)r'-C'}$ tuples, all of which are distinct due to the prefix-code constraint. Each of these tuples can now be encoded by $\mathrm{enc}(\cdot)$ to form a distinct value of k.

The total number of possible k values is therefore at least

$$N(s)^{(s-3)r'-C'} \propto N(s)^{(s-3)r'} = g^{r's} = g^r \,,$$

proving the claim. □

As can be seen, the constraints introduced by the claim regarding the values of s and of k are of much smaller effect than the constraint introduced by the proof itself: instead of looking at all of $\text{span}([s]^{c_{k'}})$, we have limited ourselves to just those k' values that can be generated from $[s]^{c_{k'}}$ through one particular method: tuple encoding, where the basic building blocks are taken from $N(s)$. Even so, we have managed to show that the number of remaining Friedmans is exponential in r and have found g to be a lower bound for the exponent.

6 Step 3: Bounding the density of Friedman numbers

Our last step is to show that even with all the restrictions introduced along the way, we still have enough remaining Friedmans to produce a density of 1. To do this, we continue in the same vein, restricting the set under discussion even further. Specifically, we will now choose just one s to work with. Amazingly, that will already be enough to reach the required density. However, for that, we must choose the value of s quite carefully.

Claim 2. *There exists a value of* s*, for which* $g = N(s)^{\frac{s-3}{s}} > 10^{L(s)}$.

Proof. There are 5 radical-free integers that are 1 digit long, 82 radical-free integers that are 2 digits long and 872 radical-free integers that are 3 digits long. Let us define the constant s to be the digit concatenation of 13 copies of each of these radical-free integers. The digit length of s, $L(s)$, is $2785 \times 13 = 36205$.

The specific s we have chosen is composed of 13 copies of each of the 959 radical-free integers below 1000. If we partition s into these radical-free integers and form all tuples of length 959 that can be created by distinct permutations of these numbers, then we have $G = (959 \times 13)!/13!^{959}$ distinct tuples. This number is a lower bound for $N(s)$, because the tuples, being of constant length, necessarily form a prefix code.

Either direct calculation or use of Stirling's formula can be used to show that G is a value with 36258 digits, whereas $10^{L(s)\frac{s}{s-3}}$ has only $L(s) + 1 = 36206$ digits. Therefore, $G > 10^{L(s)\frac{s}{s-3}}$, and $g \geq G^{\frac{s-3}{s}} \Rightarrow g > 10^{L(s)}$. □

As a side note, for this particular s it is possible to choose $c_0 = c_1 = c_{10} = c_{L(s)} = 1$, and even more complicated expressions such as $10^{L(s)} - 1$ are still within the span of a single copy of s. This should come as no surprise, because s contains over 2000 copies of each of the ten digits. However, the present proof does not rely on the specific value of any c_i.

Armed with the claim, we can now prove Theorem 1 as follows:

Proof. In order to bound the density of Friedman numbers within the integers, consider all (m, k)-pairs constructable using a specific infix length $l = L(s)r$, where $k = k'l$.

Let x be an integer chosen randomly and uniformly in $[0, 10^M)$, for some sufficiently large M. We wish to bound from above the probability that x is *not* any of the Friedman numbers corresponding to any of the (m, k)-pairs in the set.

For any given (m, k), this probability is $1-10^{-l}$, because exactly l digits are restricted to a specific value. However, the probabilities relating to any two (m, k)-pairs with the same l are independent, because the restricted digits do not overlap (hence our requirement that k be divisible by l). This means that the total probability is bounded by

$$(1 - 10^{-l})^{C_0 g^r}, \tag{4}$$

where C_0 is the multiplicative constant in the $\Omega(g^r)$ bound from Claim 1. Our goal is to prove that this probability drops to zero as r tends to infinity.

When r tends to infinity, so does l, so the expression $(1-10^{-l})^{10^l}$ tends to $1/\mathrm{e}$, where e is Euler's constant. By taking the natural logarithm of the reciprocal of the expression in Equation (4), we get $10^{-l} C_0 g^r = C_0 \left(g/10^{L(s)}\right)^r$. We therefore need, equivalently, to prove that this latter expression tends to infinity with r. However, Claim 2 already showed that $g > 10^{L(s)}$, so we are done.

This proves that as n increases, the probability of an integer uniformly sampled in $[1, n]$ to be a Friedman number approaches 1, which is what we set out to prove. □

7 But. . . is any of this math?

Having read this proof, one may wonder again about the quote by G.H. Hardy from which we started, and which dismisses the study of narcissistic numbers offhand as something "suitable for puzzle columns and likely to amuse amateurs", but not a pursuit for serious mathematicians. Where does Hardy's opposition come from, and is it applicable in this case?

Well, perhaps the most obvious and most serious objection to the study of narcissism within the realm of number theory is that narcissistic numbers are defined in terms of their own digits. Because of this, the definition of a Friedman number is dependent on the base of representation. What is a Friedman number in decimal notation may not be a Friedman number in binary, and vice versa.

However, Gödel numbering can be said to exhibit similar properties: different representation models can be used to encode tuples (or claims, or computer programs) as integers. A quine program in a specific programming language is not necessarily a quine in another programming language. The important fact, from a theoretical perspective, is that any general-purpose computer language allows the construction of a quine.

In light of this, it may be said that the true test of the number-theoretic nature of this type of study is in whether the mechanism of the proof and the ultimate conclusions are agnostic to the representation model, rather than whether the representation model itself is dependent on external factors. A representation-agnostic derivation can be said to describe true number-theoretic properties, whereas one that is based on representation can justifiably be referred to as "just a game with numbers".

Going once again over the proof above, one can see that most of it made no use whatsoever of the specific representation used. Only one part, namely the proof of Claim 2, seems to be truly dependent of the base of representation, and even it only requires that the number of radical-free numbers of certain digit lengths will be large enough. Compare this with an argument such as the one given in Equation (1).

In [1], it was shown that even this slight dependence on the base of representation can be omitted. In every base of representation, there is an s value, generated in much the same way as was done here, that satisfies the equivalent of Claim 2 for that base of representation, and therefore proves a density of 1 for the base-b Friedman numbers in any base b.

To me, this seems like a convincing-enough argument that what is shown here is not just mathematical techniques coming to the aid of a non-mathematical problem, but rather true math. Would it have convinced Hardy? If I ever see him, I'll ask.

8 Next stop: the nice ones and the vampires

We showed how the density of Friedman numbers can be computed in any base of representation and have proven it to be 1. One can now continue to ask similar questions regarding the density of several commonly-mentioned subsets of Friedman numbers, which are interesting in their own rights. Among these are vampire numbers [9], first introduced by Clifford A. Pickover, and "nice" Friedman numbers [11], introduced by Mike Reid. Both sets restrict in some meaningful way the computation generating the numbers in the set. A vampire number utilizes no exponentiation, whereas a nice Friedman number requires the computation to retain the number's original digit order.

In [2], it is shown that the infix method presented here can be made to work on nice Friedman numbers (and prove a density of 1), but only for the binary, ternary and quaternary nice Friedmans. It is not known whether the same method also works for higher bases and, specifically, for the decimals (although, for bases of representation higher than 27 the method is known to break down). For vampire numbers, the infix method breaks down entirely, regardless of the representation basis.

Of course, a method that breaks down, just like a method not provably known to work, is neither a proof that the density is 1 nor a proof of the opposite. The density of neither nice Friedman numbers nor Vampire numbers is presently known.

References

[1] M. Brand. Friedman numbers have density 1. *Discrete Applied Mathematics*, 161(16–17):2389–2395, 2013.

[2] M. Brand. On the density of nice Friedmans. http://arxiv.org/abs/1310.2390, Oct. 2013. [Online; accessed 10 October 2013].

[3] P. Bratley and J. Millo. Computer recreations: Self-reproducing automata. *Software – Practice & Experience*, 2:397–400, 1972.

[4] E. Friedman. A036057. http://www.research.att.com/~njas/sequences/A036057. [Online; accessed 23 August 2010].

[5] E. Friedman. Problem of the month. http://www2.stetson.edu/~efriedma/mathmagic/0800.html, Aug. 2000. [Online; accessed 23 August 2010].

[6] K. Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme, I. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931.

[7] G. Hardy. *A Mathematician's Apology*. Cambridge University Press, 1940.

[8] R. Muller. A007916. http://www.research.att.com/~njas/sequences/A007916, Mar. 2010. [Online; accessed 6 October 2013].

[9] C. Pickover. *Keys to Infinity*, chapter Vampire Numbers, pages 227–232. Wiley, New York, 1995.

[10] W. Quine. *The Ways of Paradox and Other Essays*, chapter The Ways of Paradox, pages 1–21. Harvard University Press, Cambridge, 1966.

[11] D. Rattner. A080035. http://www.research.att.com/~njas/sequences/A080035, Mar. 2003. [Online; accessed 23 August 2010].

[12] C. Rose. Radical narcissistic numbers. *Journal of Recreational Mathematics*, 63:250–254, 2004–2005.

[13] A. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(2):230–265, 1937.