Paideia Real Numbers, Continued:

Real numbers that are not algebraic are called "transcendental". We have just proved (in part 1) that transcendental numbers must exist (since algebraic real numbers are countable but all real numbers are uncountable), indeed an uncountable infinity of them, but we have not produced a single one. Spivak gives a nice proof in his calculus book that the famous number "e", the base of the natural logarithms, is transcendental. It turns out also that π, the circumference of a circle of radius 1/2, is transcendental. So these are very exotic numbers. $2^{1/2}$ is an unusual number but if you square it you get an integer so it really isn't too strange. But no matter how large a power of π you take you never get an integer. Also if you add up a whole lot of different powers of π you still don't get an integer. (Why not?) Since transcendental numbers are so different from the more familiar numbers it is hard to study them, hard to think of examples or how to name them, and hard to recognize whether a given number is one of them. Which means basically that I don't happen to know anything about them, so I'm telling you that they are hard to understand. How would I know? I never tried. So ignore what I said.

Yes, there really is a square root of 2

Now that we have proven that there is no rational number whose square is 2, and have introduced the much larger class of numbers represented by all infinite decimals, we should justify our construction by proving that there does exist an infinite decimal whose square is 2. Recall we don't even know yet how to add and multiply infinite decimals. So this will take us a little while. We will eventually follow the same approach you already are familiar with which lies behind our use of 1.414 to approximate the square root of two. That is, $(1.414)^2 < 2 < (1.415)^2$. So we use as an approximation to the square root of 2, to the nearest n^{th} decimal place, the largest such number whose square is less than 2. So we must begin by discussing the notion of "order" on the real line, which merely means determining when one number is larger than another.

Discussions of how large, or equivalently how small, a number is have great importance for us, since a real number (an infinite decimal) is essentially just something that can be approximated by certain rational numbers (i.e. by the associated sequence of finite decimals). Thus it is interesting to know how close a given rational number is to the real number it is supposed to approximate, or equivalently, how small is their difference.

[From now on, virtually everything in analysis, every problem solved throughout calculus and beyond,(as opposed to those solved by finite algebra), will always proceed to specify the solution to a particular problem simply by giving an infinite sequence of approximate solutions to the problem, and showing that the actual solution is the unique number, or object, which is being approximated by that sequence. For example, to prove that there is a real number whose square is equal to 2, we produce an infinite sequence of rational numbers whose squares are none of them equal to 2, but whose squares get closer and closer to 2, and we say the real square root of 2 is the real number determined by our sequence of rationals. Of course we have to do that somewhat carefully. And before we can even begin to attempt this we have to know how to square a real number.]

1) Remark: Before we start trying to define addition and multiplication for infinite decimals let's assume that we know how to do these things for finite decimals, since in fact all of us do know this. Let us assume also that all the familiar properties are known to hold for these operations, which means the properties Spivak calls P1 - P12, in chapter 1 of his book. In particular we know what it means for one finite decimal to be greater than another and we know therefore what it means to be positive (i.e. greater than 0). We will write "$a > b$" for "a is greater than b" and "$a < b$" for "a is less than b". Then we also assume, for finite decimals, that if $a < b$ and $c > 0$, then $ac < bc$. Also if $a < b$ and c is any number then $a+c < b+c$. We know also that .000000....1, with a 1 in the n^{th} place, is equal to 1 divided by 10^n. Now we try to extend these properties, so that they make sense, and are still true, for the case of infinite decimals.

Ordering, upper bounds, and least upper bounds

2).Remark: Because I wished to motivate the concept of a real number by evoking its historical origin as a tool for measuring, i.e. comparing, lengths we began from the notion of points on the line, for which we had to invoke certain axioms, the Archimedean and completeness axioms. Now that our intuitive motivation is over you may forget all of it if you wish and just start from scratch now, considering that a real number is an infinite decimal, but remembering that those represented by finite decimals, and only those, have two decimal representations, the second one ending in all nines.

3). Definition of "greater than":
Given two (positive) infinite decimals how do you tell which one is bigger? First make sure you don't have one that ends in all nines; [if you do, change it for the other one, the finite version that ends in all 0's.] Okay, then the one with the bigger integer part, if there is one, is the bigger number. Or, if they both have the same integer part, then the one with the bigger tenths digit is the bigger. If they have the same tenths digit, then the one with the bigger hundredths digit is the bigger. And so on. I.e. unless the two numbers are the same, then eventually one of them has to have a bigger digit, and as soon as one does, then that one is the bigger number regardless of what the rest of the digits are.

4). Exercise: Define how to tell which of two infinite decimals is bigger, when one or both of them may be negative. Check that when applied to finite decimals our definition of "greater than" has the right meaning.
5). Exercise: Prove that if x and y are any two real numbers, that precisely one, and not more, of these statements is true: x equals y, or x is less than y, or x is greater than y .

Now we can **prove** the Archimedean properties:

6). Theorem:(Arch I): If x is any real number then there is an integer n which is bigger than x.
Proof: Just take n to be one more than the integer part of x.

7). Theorem:(Arch.II): If x is any real number bigger than zero, there is an integer n so that 1/n is smaller than x.
Proof: If the integer part of x is bigger than zero, then $.1=1/10$ is smaller than x. That takes care of that case. If the integer part of x is zero, just go out until you find a non-zero digit in the expansion of x. Then the number which has all zeros one place farther out than x does, but

then has a 1, and all zeros after that, is smaller than x and has the form $1/(10)^k = 1/n$, where $n=(10)^k$.

We get a very important corollary from this seemingly simple property:

8). Corollary: If x is a non-negative number which is smaller than 1/n, for every n, then x is zero.
Proof: This is a rephrasing of the previous theorem, which says if it were not zero then some number of form 1/n would be smaller than it. **QED.**

Actually since I've been so picky about our not knowing how to subtract or add yet, we need the slightly stronger version of Archimedes' axiom:

9).Lemma (Arch.3): If x and y are positive real numbers and x<y, then there is a finite decimal a that lies between them in the sense that it is greater than x and less than y. i.e. such that
$x < a < y$.
Proof: First make sure neither x nor y ends in all 9's. Then, suppose x and y agree out through the n^{th} decimal place but y is larger in the $(n+1)^{st}$ place. Since we may assume x does not end in all 9's, it is possible to go out further than the $n+1^{st}$ place and find an entry in x which is less than a 9. Let a be the finite decimal obtained by letting a agree with x out to that digit but then replace that digit by the next larger digit. Then complete a by putting all zeros after that. Then a is greater than x but less than y. **QED**.

The class gave a beautiful solution of the problem of how to define the square of two positive real numbers, as reported on by Nicos, in Athens on Friday. From now on, this is an attempt to write up your solution, and to include the refinements that we began on Friday, i.e. the definition of "approaching zero", and the proof that, as you well understood, this is indeed the necessary extra ingredient to insure that a nested sequence of intervals will have only one common point.

Definition: A sequence $\{r_n\}_{n\geq 1}$, of real numbers is said to approach zero, as n approaches infinity, if and only if, for any given positive real number ∂ there exists a corresponding positive integer N, such that for every $m\geq N$, we have $|r_m|\leq \partial$. In shorthand, we write r_n-->0, as n-->infinity.
[That is, if you tell me how small you require the numbers to be, say less than some positive number e, then it must be true that if you go out far enough in the sequence, for example farther than the N^{th} entry, then all the numbers which are that far or farther out, are guaranteed to be smaller, in absolute value, than your ∂. This says that although I do not claim that the numbers ever get to be equal to zero, they do eventually get as close to zero as you want. The question of finding a limit is not one of determining what number is actually reached, but of determining what number is being approximated.]

Here now is the basic uniqueness result:

10) Corollary: If x and y are <u>any</u> two real numbers, and if $a_0, a_1, a_2, \ldots,$ and $b_0, b_1, b_2, \ldots,$ are two infinite sequences of <u>finite</u> decimals such that we have $a_n \leq x \leq y \leq b_n$,for all n , and such that (b_n-a_n) approaches zero, as n approaches infinity, then x=y.
Proof: If not, so that in fact $x < y$, then the previous corollary would give us a finite decimal d such that $x < d < y$. Applying the corollary again gives us another finite decimal e such that $d < e < y$. Then we would have $a_n < d < e < b_n$ for all n, where now all these numbers are finite decimals. Subtracting d from the last three of them now gives $0 < e-d < b_n-d$. On the other hand subtracting a_n from the first two of them gives $0 < d-a_n$, and adding b_n-d to this gives $b_n-d < b_n-d+(d-a_n) = b_n-a_n$. Putting these all together gives $0 < e-d < b_n-d < b_n-a_n$, for all n. This, however gives a contradiction, since then $\partial =(e-d)$ would be a positive number that (b_n-a_n) always stays greater than, in contradiction to the assumption that (b_n-a_n) "approaches (arbitrarily near to) zero". **QED**.

11): Corollary: If $[a_1,b_1], [a_2,b_2],\ldots$ is an infinite sequence of closed bounded nested intervals, whose end points are finite decimals, and whose lengths, (b_n-a_n), approach zero, then there is at most one real number which is in all of the intervals.
Proof: This is essentially what was proven.

Now we will begin the proof of the existence of at <u>least</u> one such point, i.e. of the completeness property.

<u>12). Upper bounds</u>: Let S be a set of real numbers, and call b an upper bound for S if b is a real number and if b is at least as big as every number belonging to S. (Note that if S is the empty set then every number is an upper bound for S.)

<u>13). Least upper bounds</u>: If S is a set of real numbers, a least upper bound for S is a number L such that L is an upper bound for S, and no number smaller than L is an upper bound. [Note that it is not required that the least upper bound should itself belong to the set S, but only that it shall be a real number.]
Advice: Memorize those important definitions exactly as they are given.

14). Problem: Prove that the empty set has no least upper bound, and that any set has <u>at most one</u> least upper bound.

The following theorem expresses the most absolutely fundamental property of the real numbers which distinguishes them from the rationals:

<u>15). Theorem: Every non empty set S of real numbers which has an upper bound, has a least upper bound</u>.
Proof: Again assume our numbers are all positive (to make them easier for me to think about). We think of the reals as given by infinite decimals, as usual. Note first that since our set S is bounded above, some integer is an upper bound, by Arch. 1. Moreover since the set is non-empty, some integer fails to be an upper bound. Thus there is a smallest integer which is an upper bound. This is not of course necessarily the (real) least upper bound we are looking for , since it may not be the least real number which is an upper bound, it is merely the least <u>integer</u>

which is one. Now we define the least (real) upper bound x of the set S one digit at a time. To define the integer part, take the largest integer that is not an upper bound for S, i.e. one less than the smallest integer upper bound. Now among the numbers { n.0, n.1, n.2, n.3, n.4, ..., n.9}, there is exactly one such that it is itself not an upper bound for S, but such that you get an upper bound by adding .1 to it. Let a_1 be the digit between 0 and 9 in which this number ends. Then we state that our number x starts out $n.a_1$. Now among the numbers $\{n.a_10$, $n.a_11$, $n.a_22$, ..., $n.a_19\}$ there is again exactly one such that, it is not itself an upper bound for S, but such that you get an upper bound by adding .01 to it. Let a_2 be the appropriate digit, and then our number x starts out as $n.a_1a_2$. Continuing in this way we get at least a prescription for constructing an infinite decimal $x=n.a_1a_2a_3$........., and one which I claim is the least (real) upper bound of S.
We must check two things:
1) x is an upper bound for S
2) No number smaller than x is an upper bound for S.
Let's try 2) first. Now our construction may yield for x a decimal that ends in all 9's, (apply it for instance to the set S={.9, .99, .999, ...}), but even if it does, any number smaller than x will be given by a decimal which equals x up to some point and then has a digit which is smaller than the corresponding digit of x. So if y is a smaller number than x, then look at the first digit y has which is smaller than the corresponding digit of x. By construction of x, the decimal which agrees with x out to, and including, this digit is not an upper bound for S. Moreover it is at least as large as y, so that y too is not an upper bound. QED for 2).
Now we try 1). If x were not an upper bound for S, then there would exist a number z in S which is larger than x, i.e. x<z and z´S. Now if $x=a_0.a_1a_2$..., then write $x_n=a_0a_1a_2...a_n$, for the finite decimal that agrees with x out through the n^{th} decimal place and then has all zeros. Also, define $y_n=x_n+(10)^{-n}$, so that (by construction of x), for all n, x_n is not an upper bound of S but y_n is an upper bound. Since zŸS then $z\leq y_n$ for all n. Thus we have
$x_n \leq x < z \leq y_n$, for all n. Since however $y_n-x_n = (10)^{-n}$ ¨ 0, this contradicts Theorem 10, which says that then x = z. **QED**.

Now we can prove the completeness property for infinite decimals which, for points on the line, we earlier took as an axiom.

16). Theorem: **(Completeness of the real numbers)**: If $[a_1,b_1]$, $[a_2,b_2]$,.....is any infinite sequence of closed, bounded, nested intervals, whose endpoints are finite decimals, and such that the sequence of their lengths, (b_n-a_n), approaches zero, then there is exactly one real number which lies in all the intervals.
Proof: Corollary 11 shows there is at most one such point so we have to prove there is at least one. Let S be the set of all left end points of the intervals, i.e. $S=\{a_n\}$. Then b_1 is an upper bound for S and a_1 is in S so S is not empty and is bounded above, and thus has a least upper bound, which we call x. Then we claim that x lies in all the intervals. Since x is an upper bound for the a's, we have $a_n\leq x$, for all n. Moreover, since every b_m is an upper bound for the set of a's, and x is the least upper bound no b can be less than x. I.e. $x\leq b_n$ for all n. Thus for all n, $a_n\leq x\leq b_n$, i.e. x is in all the intervals.QED.

17). Lemma: If c is any finite decimal, then the sequence $\{c/(10)^m\}$ approaches zero as m approaches infinity.
Proof: Let $\partial>0$ be any given positive number. Then by Theorem 7 we can find an n such that $(1/n)<\partial$ Now by Theorem 6 choose an integer k so large that $k>c$. Then since $(10)^k > k > c$, we have $(c/(10)^k) < 1$. Moreover, $(10)^n > n$, so that $1/(10)^n < 1/n$. Thus $(c/(10)^{k+n}) < 1/n$, too. Do you see why? Thus for $N=n+k$, we have $(c/(10)^m) < \partial$, whenever $m \geq N$. Consequently, the lemma is proved.

How to add two positive real numbers:
18). Definition/Theorem:
Let $a=a_0.a_1a_2a_3a_4....$, and $b=b_0.b_1b_2b_3b_4.......$, be the two reals. To add them, Steve's idea was to go back to the sequence of intervals with end points given by finite decimals, add the endpoints, and get a new sequence of intervals and then try to show this new sequence of intervals contained a unique point. That point of course should be a+b. So recall the appropriate sequences of intervals: from the decimal for a, we have $\{ [a_0,a_0+1]$, $[a_0.a_1,a_0.a_1+.1]$, $[a_0.a_1a_2,a_0.a_1a_2+.01]$,..........., and then from b, we get the same sort of sequence only with b's instead of a's. Okay, now Steve's construction tells us to add the left endpoints of the a-intervals to the left endpoints of the b-intervals in order to get the left endpoints of a new set of intervals, then do the same with the right endpoints. So we have these intervals: $[a_0+b_0,a_0+b_0+2]$, $[a_0.a_1+b_0.b_1,a_0.a_1+b_0.b_1+.2]$, $[a_0.a_1a_2+b_0.b_1b_2,a_0.a_1a_2+b_0.b_1b_2+.02]$,........ Now we must check that these intervals do indeed have exactly one common point. Since all the intervals are closed and bounded by definition, we must check that they are nested. Looking at the sequence of left endpoints, we see that they are getting larger, or at least not getting smaller, and if you look hard at the sequence of right endpoints you will see that they too are either getting smaller or at least not getting larger. To say it another way, the intervals <u>are</u> nested. Now another look at these intervals shows that, for the n^{th} interval, the right endpoint minus the left endpoint is equal to $2/(10)^n$. Consequently, (b_n-a_n) ¨0, so that, by thm. 16, indeed the intervals contain exactly one common point. **QED**.

Definition: The unique point common to all the intervals in definition 11 is defined to be the sum, a+b, of the real numbers a and b.

Remark: After Steve's insight has shown us the way to define this, we can look back with hindsight and observe that certainly this was the right way to do things. Why? Well we wanted to define addition so that the same properties continued to hold which held for adding rational numbers. That means that if $a\leq x\leq b$, and $c\leq y\leq d$, then we would want to have $a+b\leq x+y\leq c+d$. Steve's definition is precisely the one which is forced on us if we want this property to continue to hold. I.e. it is the only possible definition which has a chance of making this property continue to be true for all reals. That is , either this is the good definition, or else there is no good definition. To check that it really is the good definition, we would have to prove that all the usual properties do hold for this definition, i.e. Spivak's properties P1-P12.

19). Problem: Note that we have been assuming our numbers were positive. How do you define addition of a positive and a negative number? I.e. how do you define subtraction?
20). Exercise: Use this definition to check directly that a+b=b+a, for any two real numbers.
21). Exercise: Think about, and write down, what would be the first few steps of checking associativity.

Now we recap Nicos's explanation of how to define the square of two real numbers.

How to define squaring, as solved by the class
Let $x=a_0.a_1a_2...$ be any real number, and define for each n, $x_n=a_0.a_1a_2..a_n$, and $y_n=a_0.a_1a_2...a_n+(10)^{-n}$. Then $x_n \le x \le y_n$, for all n. We consider the intervals $[(x_1)^2,(y_1)^2]$, $[(x_2)^2,(y_2)^2]$, ..., and the claim is that they contain a unique point which we will then define as x^2. Since we are assuming the familiar properties of squaring for the finite decimals, it follows that this is a nested sequence of bounded closed intervals, so one must only check that their lengths approach zero. We need a definition and a lemma:

Definition: A sequence $\{r_n\}$ is called bounded if there is some real number B such that $|r_n| \le B$, for all n. I.e. there is some one number which is at least as big as the absolute value of everything in the sequence.

22). Lemma: If $\{r_n\}$ is a sequence of finite decimals which approaches zero, and if $\{s_n\}$ is a bounded sequence of finite decimals , then the sequence of products $\{r_n s_n\}$ also approaches zero.
Proof: Let $\partial>0$ be any given positive number. We must show it is possible to pick an integer N so that if $m \ge N$, then $|r_m s_m|<\partial$. Pick an integer k so that $(1/k)<\partial$ Then consider $\partial_1= 1/(kB)$, where B is an integer and a bound for $\{s_n\}$. Since $\{r_n\}\text{-->}0$, there is an integer N such that for $m \ge N$, we have $|r_m|<\partial_1=1/(kB)$. Then $B|r_m|<1/k$. Moreover, since B is at least as big as every $|s_n|$, $m \ge N$ also implies that $|r_m s_m| \le B|r_m| < (1/k) < \partial$. **QED**.

Corollary: The lengths of the intervals $[(x_n)^2,(y_n)^2]$ considered above, do indeed approach zero.
Proof: The lengths are equal to $(y_n)^2-(x_n)^2 = (y_n- x_n)(y_n + x_n)$. Since $(y_n-x_n) = (10)^{-n}$, the sequence of first factors approaches zero. Since $2x_1 \le y_n+x_n \le 2y_1$, for all n, (why?), the sequence of second factors is bounded. Consequently the product, i.e., the sequence of lengths, goes to zero as claimed. **QED**.

Corollary: If x is any real number and if x_n and y_n are defined as above then there is a unique real number which we will call x^2, such that for all n, we have $(x_n)^2 \le x^2 \le (y_n)^2$.
Remark: From the proof of the completeness property it follows that we could have equivalently defined x^2 as the least upper bound of the sequence $\{(x_n)^2\}$.
Exercise: If x and y are two positive reals such that $x < y$, then $x^2 < y^2$.

At last we can prove there is a square root of two.

Theorem: There does exist a real number x such that $x^2=2$.

Proof: We define x one digit at a time. Let a_0 be the largest integer whose square is not greater than 2, i.e. $a_0=1$. Then let a_1 be the largest digit from 0 through 9 such that $a_0.a_1$ has square less than 2, i.e. $a_1=4$. Continuing in this way define $x=a_0.a_1a_2$....... so that for each n, $x_n=a_0.a_1a_2...a_n$ is the largest finite decimal which has all zeros after the n^{th} decimal place, and whose square is less than 2. Then if $y_n=x_n+(10)^{-n}$, we have for all n, $(x_n)^2<2<(y_n)^2$. But then the real number 2 satisfies the property that defined the number x^2, according to our definition~ That is, it lies in all the intervals defining x^2, and since no other number does, we must have $x^2 = 2$.

#22). Problem: Use Steve's method to define the product of two positive real numbers.

#23). Problem: Prove that a positive real number has a multiplicative inverse. (Given a real number x >0, tell how to define a certain suitable real number y, and then <u>prove</u> that xy=1).

Remark: Without going through all the details, I propose now to take for granted that our new definitions for adding and subtracting and multiplying and dividing real numbers can be made also for negative reals and have all the properties P1-P13, (P13 is in chapter 8 of Spivak). [As usual, we define division as multiplication by the inverse.] Those are all the properties needed to completely describe the real numbers and to distinguish them from all other sets of numbers. The proof that any set of numbers that also has these properties must be exactly like (i.e. "isomorphic to") the real numbers is given in chapter 28 of Spivak. You may find it interesting. It is not as tedious as the actual construction of the real numbers, of which we have done only a portion.

Factoring integers into primes

Now let's pause and do some arithmetic properties of the integers.
A fundamental property of the positive integers, which is a powerful tool for proving things about them is the following "well-ordering" property:

Axiom: Every non-empty subset S of the positive integers contains a smallest element.
[**Remark**: Since this element is required to actually be in the set S, this is different from the least upper bound property. This new property is false for the positive reals. For instance the infinite set of reals of form 1/2, 1/3, 1/4, 1/5, ..., contains no least element. Its greatest lower bound is zero, but that is not in the set.]

Theorem: Every positive integer greater than one can be factored into positive prime factors in exactly one way.

Proof: The existence is the easy part. If there is an integer x>1 that does not factor into primes, then by the well-ordering axiom, there is a smallest such positive integer. I.e. there is an x>1 that does not itself factor, but such that every smaller integer greater than one does factor. But this is impossible, for our number x cannot itself be prime, or else it would already be factored into primes, with just one factor! Therefore it can be factored into two factors both of which are smaller than x and greater than 1. But then both of the factors, being smaller than x, <u>will</u> have prime factorizations, and then by putting the two factorizations together, we would

get a factorization of x. This contradiction shows that no such smallest nonfactorable number x can exist, and hence indeed no non-factorable numbers exist at all.
Uniqueness is harder, and may have been proved first by Gauss. The key lemma usually used nowadays, is to show: if a prime integer p divides a product of positive integers ab, then p divides either a or b. A nice way to do this follows from a classic fact about "greatest common divisors" proved in Euclid. It can be stated as a fact about measuring lengths using two different rulers, which are commensurable, i.e. whose ratio of lengths is a rational number. The basic result is that the shortest length one can measure by using both rulers, equals the longest length that can itself be used to measure both rulers. I.e. given two integers a,b, the smallest positive integer that can be written in the form an+bm where n,m are any integers, either positive or negative, equals the largest integer d such that d divides both a and b evenly. Assuming this, if p is prime and divides ab but does not divide a, then the largest integer that divides both p and a is 1, hence 1 can be written in the form 1 = an+pm for some integers n,m. Then multiplying by b gives us b = abn + bpm. Now assuming p divides the product ab, it follows that p divides both terms on the right side of the equation, hence divides also the left side. Thus we have shown that if p divides ab but does not divide a, then p divides b.

Now assume we have factores some number n into primes in two ways n = p1p2---pr = q1q2....qs. Since p1 is prime and divides the left side it also divides the right, hence p1 divides some q, which we may renumber as q1. But q1 is prime so if p1 divides it, since p1 ≠ 1, it must be that p1=q1. Then we can cancel p1 and q1 on both sides and have a new equation p2p3...pr = q2q3...qs. We now apply th same argument to the prime p2, eventually canceling it with some prime q we may renumber as q2. Eventually we have canceled all primes on both sides, in particular, each prime pj was equal to some prime qk, and the factorization was unique. **Q.E.D**.

We can prove now that the real number whose square is two is not rational. I.e. if n/m is a rational number whose square is 2, then n^2/m^2 = 2, so n^2 = 2m^2. But in the prime factorization of n^2 there are twice as many 2's as in the factorization of n, hence an even number, and the same holds for m^2. But that means in the oprime factorization of n^2 there occur an even number of factors of 2, while in the factorization of 2m^2 there occur an odd number (the even number of factors of 2 in m^2, plus the extra "2" in front of 2m^2). Since an integer has only one prime factorization, it cannot be true that n^2 = 2m^2, so 2 ≠ n^2/m^2.