Patterns, Rules, and Inferences

Achille C. Varzi (Columbia University, New York)

[Final version published in Jonathan E. Adler and Lance J. Rips (eds.) *Reasoning: Studies of Human Inference and Its Foundations*, Cambridge: Cambridge University Press, 2008, pp. 282–290.]

1. Introduction

Some of our reasoning is strictly *deductive*; we conclude that the available evidence supports a certain claim as a matter of logical necessity. For example, the following reasoning is deductive, since it is not possible for the conclusion to be false if both premises are true.

(1) Every F is G, and x is F. Therefore, x is G.

Often, however, we are not in a position to produce a deductive argument; often we can only establish that the evidence supports the conclusion to a high degree of probability. Such *inductive* reasoning, as it is normally called, is in turn divisible into two types, according to whether or not it presupposes that the universe or some relevant aspect of it is law-like, or rule-governed. Reasoning that does not require this presupposition may be classified as *statistical*, since the evidence described by the premises supports the conclusion for purely mathematical reasons. For example, the following inductive reasoning is statistical:

(2) Almost every F is G, and x is F. Therefore, x is G.

Here it is rational to reach the conclusion even though it does not follow as a matter of logical necessity, for the probability of x's being G is, given the facts, much higher than the probability of x's not being G (other things being equal). The second type of inductive reasoning is generally classified as *Humean*, after the philosopher who first studied it thoroughly, and corresponds to those arguments that do require the presupposition of law-likeness. The following is an example:

(3) Every F previously observed was G, and x is F. Therefore, x is G.

Again, the conclusion does not follow as a matter of logical necessity, so the argument is nor deductive. Yet the available evidence gives excellent reasons to believe in the conclusion rather than in its negation. Unless the relationship between being F and being G is random, the evidence strongly suggests the existence a law to the effect that *every* F is G.

Humean arguments are of great practical utility, since we often need to reach conclusions and make decisions on the basis of evidence that is neither conclusive (thus preventing us from reasoning deductively) nor complete (preventing us from reasoning statistically). The presupposition of law-likeness, however, plays a crucial and controversial role, and figuring out exactly what role it plays is no straightforward business.

2. The Game of the Rule

Consider the following familiar game. X thinks of a certain sequence (say, a sequence of numbers) and Y must figure out what the sequence is. To get started, X gives an initial fragment of the sequence. Y must look at it carefully and, on the basis of what she sees, she must try and figure out how the sequence continues. Which is to say: she must figure out the underlying pattern, uncover the rule by means of which the sequence is generated.

For example, let us focus on (infinite) number sequences. If X's initial segment looks like this:

(A) 1, 3, 5, 7, 9, 11, …

then Y is likely to come up with a quick and reasonable guess: The sequence must consist of the positive odd integers, in their natural order. If the initial segment looks like this:

(B) 1, 2, 3, 5, 7, 11, …

then, again, Y may easily figure out how to continue—hence the rule by means of which the sequence is generated: This is the ordered sequence of the prime numbers, i.e., the positive integers that do not have any other integer factors except for 1 and themselves. If Y guesses the rule within the allotted time, she wins the game. Otherwise X wins.

Now, some cases are more challenging than others, of course, and this is where the game gets interesting. For instance, consider the segment

(C) 1, 3, 6, 10, 15, 21, 28, …

This is the beginning of the sequence of the so-called triangular numbers, namely, those numbers that equal the sum of consecutive integers beginning with 1. More precisely, the rule underlying this sequence is that the *n*th element, S_n, is the sum of the first *n* positive integers:

(C*) $S_n = 1 + 2 + 3 + \ldots + n$.

(For example, the fourth triangular number is 10 = 1 + 2 + 3 + 4.) One may want to run a test and see how people actually perform, but a good guess would be that in this case it takes more thinking to figure out the solution. Y might even object that she has never heard of the triangular numbers—whereas she had heard of even and prime integers—so how could she figure out the rule? Still, X may just answer that one need not *know* what a triangular number is in order to see the pattern. With some patience, Y could still figure out that the sequence obeys the rule defined in (C^*). Or she could figure out the rule under a different, more intuitive description. For example, Y might realize that the numbers in the sequence correspond to the different ways in which we can form a triangular array of dots, or bowling pins, or billiard balls, as in the following diagram:

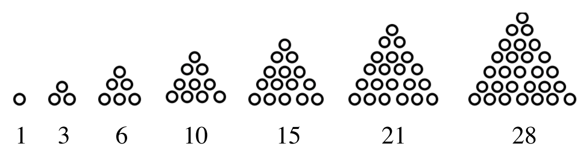

So Y could describe the sequence in terms of this intuition:

(C^{**}) S_{n+1} = the smallest number of dots (pins, balls) that are needed to form a triangular array of size greater than S_n, starting from $S_1 = 1$.

(This is actually why these numbers are called *triangular*, in analogy with the *square* numbers, which correspond to the different ways in which we can form a square array of dots.) How one comes up with the rule and how one describes it—X may insist—is not important in order to win the game. It is only important in a derivative sense, namely insofar as it makes the game playable by people with different backgrounds. The game is interesting precisely because the mental process whereby the rule is uncovered may involve different sorts of cognitive insight.

In fact, it is worth noting that although in these examples the rule by means of which the sequence is generated is essentially determined by its number-theoretic properties, it need not be so in general, even if the sequence consists of numbers. For instance, suppose X offers the following segment:

(D) 1, 22, 333, 4444, 55555, …

This is the beginning of an obvious sequence and, as it turns out, there is a mathematical key to this sequence, corresponding to the equation

(D^*) $$S_n = n \cdot \frac{(10^n - 1)}{9}.$$

But of course Y is more likely to describe the sequence on the basis of a different criterion, which reads the rule directly off the visual pattern exhibited by its elements:

1	One 1
22	Two 2s
333	Three 3s
4444	Four 4s
55555	Five 5s
⋮	⋮

In that case, Y's rule would not be (D^*) but rather something like this:

(D^{**}) S_n = the string consisting of the number n repeated n times.

X himself, in giving the initial segment, may have thought of the sequence in terms of (D^{**}), not (D^*), so the fact that there is a number-theoretic description of this sequence is entirely irrelevant. And in some cases there is no number-theoretic description at all, as in

(E) 1, 3, 4, 5, 7, 8, 9, 12, 14, 17, 18, …

Here X may be thinking of a rule that can be defined with reference to a linguistic property and that concerns the numerals, not the numbers:

(E^*) S_n = the nth integer whose name in English has an even number of vowels.

A favorite example of this sort is actually one that does not depend on purely linguistic considerations, just as it does not depend on purely arithmetical considerations, and is due to the American mathematician John Conway:[1]

(F) 1, 11, 21, 1211, 111221, 312211, 13112221, …

Is there a rule behind this sequence? If we look for a purely arithmetical or linguistic key, we won't find any. We must look at the sequence from a different perspective. Exactly what perspectives one may consider is of course hard to tell. But if we start reading the sequence aloud we might get a clue. Let's not read it like this:

1	One
11	Eleven
21	Twenty-one
1211	One thousand, two hundred, eleven
111221	One hundred eleven thousand, two hundred, twenty-one
312211	Three hundred twelve thousand, two hundred, eleven
13112221	Thirteen million, one hundred twelve thousand, …
⋮	⋮

[1] J. H. Conway, 'The Weird and Wonderful Chemistry of Audioactive Decay', *Eureka* 46 (1986): 5-16.

Let us read it like this:

1	1
11	One 1
21	Two 1s
1211	One 2 and one 1
111221	One 1, one 2, and two 1s
312211	Three 1s, two 2s, and one 1
13112221	One 3, one 1, two 2s, and two 1s
⋮	⋮

Then we suddenly realize what is going on: this is a "self-describing" sequence. It begins with 1 and then goes on to describe itself, in the sense that each subsequent term gives an "audioactive" description of its predecessor. The rule can be put thus, where '$d_i|r_i$' designates the string obtained by repeating r_i times the digit d_i:

(F^*) If S_n is the string $d_1|r_1 \ldots d_k|r_k$ ($d_i \neq d_{i+1}$ for all $i < k$), then S_{n+1} is the string $r_1d_1 \ldots r_kd_k$, starting with $S_1 = 1$.

3. The Rules of the Game

So much for this familiar game. It takes a moment now to realize that the game is a good model of what goes on when we engage in Humean inductive reasoning. For the game of the rule is a familiar one, not only insofar as it is often played for fun, or for pedagogical purposes (elementary school teachers often rely on it to explain—for instance—certain basic arithmetical concepts and operations); it is also familiar precisely because we find ourselves playing it all the time in our daily interactions with the world around us. We are constantly trying to figure out the rules or laws that govern the natural world, or the social world, or the stock market. We look at the facts and we try to figure out the underlying pattern, so as to predict what will happen next, exactly as we try to figure out the pattern of a sequence of numbers on the basis of an initial segment. We look at our history so far—that's the initial segment—and we try to figure out the underlying rationale, so as to know what to expect next. It may sound metaphorical, but it wouldn't be so far-fetched to claim that science as a whole is engaged in a game of this sort. Every *F* observed thus far is *G* (every number in the visible portion of the sequence is prime, for instance), so we think that there is a law-like connection between being *F* and being *G*: we think that being *F* goes hand in hand with being *G* and we conclude that the next *F* in the sequence must be *G*, too That is precisely the idea behind Humean inductive reasoning. And there are researchers who would claim that being able to reason that way—to play such a game—is a distinctive trait of rational behavior.

Douglas Hofstadter and his research group, for instance, believe that this trait is close to the core—if not *the* core—of human intelligence, and that designing computer programs capable of playing the game of the rule is the deepest and most fascinating challenge that so-called "artificial intelligence" must face.[2]

To fully appreciate the import of these claims, however, it is important that we now be explicit about a few things that we have so far been taking for granted. There are, in fact, two crucial implicit assumptions that must be satisfied in order for the game to be played correctly—two tacit Rules (with the capital 'R') that the players must observe.

The first tacit Rule is that the initial segment by means of which the sequence is introduced should provide enough information for Y to figure out the solution. For instance, with

(G) 1, 11, …

rather than (F) as the sole piece of evidence, Y would hardly come up with the Conway sequence, simply because this initial segment is compatible with many other, more plausible solutions. The sequence could in fact continue in several ways, each of which corresponds to a different solution that "fits the data" equally well insofar as the data are fixed by (G). For example, it could continue in any of the following three ways:

(G_a) 1, 11, 121, 1331, 14641, 161051, …
(G_b) 1, 11, 111, 1111, 11111, 111111, …
(G_c) 1, 11, 1, 11, 1, 11, 1, 11, 1, 11, …

and each way would correspond to a completely different rule:

(G^*_a) $S_n = 11^{(n-1)}$
(G^*_b) $S_n = 1|n$
(G^*_c) $S_n = 1$ if n is odd, and $S_n = 11$ if n is even.

There are obviously lots of possibilities, and for this reasons Y would be entitled to complain if all X gave her as a starter was just the small bit in (G). For the problem is not merely that one can come up with different answers; we have already seen that sometime the same sequence can be generated or described in accordance with more than one rule, as with (D^*) and (D^{**}). The problem is that in the present case the different answers would not be equivalent: they would describe different sequences, not the same sequence in different ways.

[2] Hofstadter's views and early results in this area are documented in his book *Fluid Concepts and Creative Analogies*, New York: Basic Books, 1995.

So this is the first tacit Rule of the game, which we can approximately formulate as follows:

R1 *The initial segment must uniquely identify the sequence.*

The interesting question, of course, is whether this Rule can be successfully implemented, or even whether it can be implemented at all. We shall come back to this question shortly. First let us mention the second tacit Rule of the game, which is equally important. This second Rule says that the sequence in question cannot be a *random* sequence. For example, it would be strange if X said that the sequence in (G) continued thus:

(G_d) 1, 11, 3, 4, 5, 10, 7, 8, 9, 9, 10, 12, 12, …

It would be strange because, on the face of it, this sequence appears to continue in a totally arbitrary fashion and there seems to be no way of subsuming it under a rule, hence no way for Y—or for anybody—to describe the sequence other than by laying out each term that compose it, one after the other. For the same reason, of course, it would be strange if Y insisted that (G_d) is on equal footing with (G_a)–(G_c). For (G_a)–(G_c) do exhibit a pattern, or so it seems, whether (G_d) does not. So, as a first approximation, the second Rule of the game can be put as follows:

R2 *The sequence must not be random, i.e., it must be rule-governed.*

In a way, R2 follows from R1. For if a sequence were random, then *no* proper initial segment could uniquely identify it. Hence, by contraposition, if there is a proper initial segment that uniquely identifies the sequence, as per R1, the sequence cannot be random. In fact, this is how randomness is often defined, at least since the pioneering work of Ray Solomonoff, Andrei Kolmogorov, and Gregory Chaitin in the mid 1960s: a sequence is random if it cannot be described more efficiently than by laying out the whole sequence itself.[3] However, as we said, R1 may not be entirely in order as it stands, so it is convenient to formulate R2 independently. And again, we shall come back shortly to the important question—whether this second Rule can be properly implemented, or taken for granted. Right now the point is just that R1 and R2 are standardly assumed to hold whenever two players engage in a game of this sort, for otherwise there is no way one can succeed in guessing the sequence.

---

[3] The seminal works are R. J. Solomonoff, 'A Formal Theory of Inductive Inference. Part I', *Information and Control* 7 (1964): 224–254; A. Kolmogorov, 'Three Approaches to the Quantitative Definition of Information', *Problems of Information Transmission* 1 (1965): 1–17; G. J. Chaitin, 'On the Length of Programs for Computing Finite Binary Sequences: Statistical Considerations', *Journal of the ACM* 16 (1969): 145–159.

4. Too Good to Be True

So, are R1 and R2 in order? Not quite, unfortunately. Let us begin with R1. On closer look, this rule turns out to be just as crucial as it is unsatifiable—upsetting as this might sound to the players of the game. Let us first illustrate this negative fact with reference to the sort of cases that we have considered so far; then we may turn to generalizations.

Consider again the example with which we began—the sequence corresponding to the segment

(A) 1, 3, 5, 7, 9, 11, …

Surely we can all see a pattern here: the odd numbers. But how do we know that this is *the* pattern? How does Y know that this is the sequence X had in mind? Y doesn't know it. She sees that every number in (A) is odd, and she sees that no odd number is missing, and since she is assuming that this initial segment uniquely characterizes the whole sequence, she concludes that the dots must be filled in by the odd numbers. That is, she concludes that the underlying rule must be this:

(A^*) $S_n = 2n–1$.

Strictly speaking, however, she doesn't *know* that this is the rule any more than she knows the truth of any generalization based on a limited amount of data. The generalization is justified precisely because she is assuming that R1 is being observed. But how can that be right? How can the initial segment by means of which the sequence is introduced provide *enough information* for anybody to figure out the solution and continue the sequence by filling in the dots accordingly? As Wittgenstein famously put it: "Whence comes the idea that the beginning of a series is a visible section of rails invisibly laid to infinity?"[4]

Suppose Y says that the sequence in question consists of the odd numbers and X says: "No, it doesn't. It consists of the odd digits repeated once, then repeated twice, then repeated three times, and so on. Here is what it would look like if I continued a little longer:

(A') 1, 3, 5, 7, 9, 11, 33, 55, 77, 99, 111, …

If you want me to be more precise, I can even spell out the rule in mathematic terms:

(A^{**}) $S_n = 2n–1 \bmod 10$, repeated $(2n \text{ div } 10) + 1$ times,

[4] L. Wittgenstein, *Philosophische Untersuchungen / Philosophical Investigations*, ed. by G. E. M. Anscombe and R. Rhees, with an Eng. trans. by G. E. M. Anscombe, Oxford: Basil Blackwell, 1953, § 218.

where mod is the function that returns the remainder of the division (of the first argument by the second) and div the function that returns the division without the remainder." Is Y entitled to complain?

In a way she is: If *that* is the rule X had in mind, then X did not comply to R1 because (A) does not amount to a uniquely identifying segment. In particular, it does not uniquely identify the sequence described in (A^{**}), for the dots can be filled in in conformity to that rule as also in conformity to the rule that Y originally suggested, (A^{*}). Of course, this means that (A) does not uniquely identify the sequence of the odd numbers, either, so Y's complaint is self-defeating. But never mind that. It is a fact that relative to (A^{**})—the rule that X had in mind and that B was supposed to figure out—the segment in (A) is not informative enough, just as the short segment in (G) would not have been informative enough to identify the rule of the Conway sequence, (F^{*}). So Y's complaint is right on the mark.

On the other hand, what is X to make of this complaint? *What* would count as an appropriate, uniquely identifying segment for the rule he had in mind? Suppose X gives the longer segment in (A') rather than (A). Would that be enough? It would not. It would be enough to rule out the hypothesis that his sequence consists of the odd integers. But many other sequences would still be compatible with that initial segment. The sequence might continue in conformity to the pattern X actually have in mind, but it might also continue according to a different pattern. For example, the sequence could consist of the perfect palindromic odds, i.e., those numbers that consist exclusively of odd digits and that are the same when written forwards or backwards. All the numbers in (A') are perfect palindromic odds. But whereas X's sequence would continue thus:

(A'_a) …, 333, 555, 777, 999, 1111, …

the rest of the sequence of the palindromic odds would contain several additional, intermediate elements:

(A'_b) …, 131, 151, 171, 191, 313, 333, …

Needless to say, even (A'_a) would be ambiguous as an input for guessing X's rule, for one may still think of different ways of continuing the series. The more we go on—the longer the initial segment is—the more the alternatives look convoluted and, in some way, "unnatural". But this is precisely the point. It is not R1 by itself that imposes a plausible constraint on the game, for R1 can *never* be satisfied: any finite segment can be continued in an infinite number of ways, just as any line-segment drawn on a sheet of paper can be extended in an infinite number of ways. Rather, the constraint comes from R1 together with the additional implicit assumption that the sequence in question must be a "natural" sequence. And sad as it might sound, it is a fact that what looks "natural" to Y may very well not coincide with what looks "natural" to X—and vice versa.

It takes a moment now to realize how important this is when it comes to playing the game for real—when the player to issue the initial segment is not just someone like us but the world itself. A sequence of observed events may suggest that a certain pattern is in place and we—playing the role of player Y—eventually come to think of that pattern as revealing a corresponding law of nature. But this is not to say that the observed events uniquely identify that law. And if the next observed event is not what we expected, we can hardly voice a complaint on the grounds that the resulting sequence looks "unnatural". We must simply admit that we were wrong, and learn to live with the possibility that out next guess will be off the mark, too. Such is the limit of our inductive practices, when they are not merely statistical but strictly Humean.[5]

Let us now look at our second meta-Rule, R2. Indeed, it might be thought that this is precisely the point where R2 enters the picture. This second Rule says that the sequence to be guessed must not be random, i.e., it must be rule-governed. And for all that has been said so far, the fact that any initial segment can be continued in many different ways does not mean that it can be continued in many *rule-governed* ways. The segment in (A) can, as also the longer segment in (A'). But perhaps a sufficiently longer segment could be provided that will admit *only one* rule-governed extension. If so, then the impasse that we have just reached in connection with R1 would dissolve as soon as we plug in R2: it would be possible to uniquely identify a sequence by means of an initial segment (a sufficiently long one) on account of the fact that all the alternative ways of continuing the sequence would qualify as random and would therefore be unacceptable by R2. In fact that is precisely the intended role of this second Rule: earlier we said that R2 follows from R1, but now we see that R1 is empty unless some further constraint is added—and R2 provides such a constraint.

Unfortunately, R2 turns out to be just as useless as R1. There are two ways of making the point. The first way goes back at least as far as Leibniz, who in the *Discourse of Metaphysics* addressed the question of whether and how one could discriminate a world in which science applies from one in which it does not.[6] Imagine—he said—that someone jots down a quantity of ink spots upon a sheet of paper helter skelter ("as do those who exercise the ridiculous art of Geomancy"). Regard-

---

[5] D. Hume, *Enquiries Concerning Human Understanding*, Sec. IV. The connection between Wittgenstein's views on rule following and Hume's skepticism has been made explicit by S. A. Kripke in his book, *Wittgenstein on Rules and Private Language*, Cambridge (MA): Harvard University Press, 1982.

[6] G. W. Leibniz, *Discourse of Metaphysics*, Section VI.

less of the particular configuration that we get, Leibniz claimed that there will always be a continuous function whose graph passes through this finite set of points, a "geometrical line whose concept shall be uniform and constant, that is, in accordance with a certain formula". As far as we can tell, the existence of such a function was purely conjectured by Leibniz, but today we know that he was absolutely right. Many good ways to construct a function that does the job are now known. For example, so-called Lagrangian polynomial interpolation will do.[7] This is the sort of function that is implemented in most computer graphics programs: we click the mouse on each spot as we go over them, and the function returns a normalized curve that connects them all—like this:

ink spots Lagrangian curve

Now, if we take it that the existence of a suitable function is an indication of the fact that the pattern is not random, then it follows that no such pattern is random. And since it is plausible to suppose that every finite sequence can be represented by a corresponding pattern of ink spots, it follows that *no* finite sequence whatsoever is random. So here is the problem: when playing the game of the rule, there is no way X can provide an initial segment that is "sufficiently long" to admit of only one rule-governed extension. *Any* finite extension of any initial segment, no matter how long, will be non-random. Which is to say that R2 does not impose any restriction of the desired sort, leaving R1 in the trashcan.

The second way of making the point is this. Suppose we rely on a more austere definition of randomness. Indeed, suppose we stick to the definition of randomness mentioned earlier, which today is widely accepted: a sequence is random if (and only if) it cannot be described more efficiently than by giving the whole sequence itself. In other words, although every sequence can be described by *some* function—as Leibniz pointed out—in some cases the function in question is too complex to do the job efficiently, and we can take that to be a sign of randomness. If we stick to this definition, then we can be assured that *there are* random sequences, so the above problem does not arise. This is obvious for infinite sequences, since the total number of such sequences is uncountable, whereas there are only countably many efficient, finite descriptions. But it is easy to prove that there are also infi-

[7] See for instance H. Jeffreys and B. S. Jeffreys, *Methods of Mathematical Physics*, Cambridge: Cambridge University Press, 1983, §9.011. (The label comes from the Italian mathematician Joseph-Louis Lagrange, who discovered the method over a century after Leibniz's conjecture.)

nitely many random sequences of finite length, at least if the language in which the sequences are coded is the same as the language available to describe them. (Regardless of the alphabet, the number of sequences consisting of *n* symbols is always greater than the number of all sequences consisting of fewer symbols, hence greater than the number of all descriptions of length less than *n*.[8]) Does this help?

Unfortunately it doesn't. The problem is not that this more austere definition of randomness depends on a notion of "efficient" description that appears to be vague. We could make that more precise. Rather, the problem is that this more austere notion of randomness turns out to be *undecidable*. That is, it can be proved that there exists no effective decision procedure (intuitively: no procedure that can be implemented as a computer program) that will always deliver a definite answer to every question of the form: Is this a random sequence?[9] Sometimes we can deliver a negative answer. The sequence of the odd numbers, for example, or even a reasonably long finite initial segment thereof, is not random because we can efficiently describe it by means of a rule such as (A^*). Ditto for the sequence of the repeated odds, the sequence of the palindromic odds, and many other sequences considered above. However, *in general* we may not be in a position to determine whether a given sequence (finite or infinite) is random. All we can say, if we cannot come up with a corresponding rule, is that the sequence is random *to the best of our knowledge*—and that is not enough. For example, at first sight the initial segment of the Conway sequence given in

(F) 1, 11, 21, 1211, 111221, 312211, 13112221, …

looks pretty random. Y might have tried to come up with an algorithm to describe it and she might have failed. So Y might have been inclined to conclude that the segment is the beginning of a random sequence, when in fact it isn't. Likewise for (G_d): we have said that the series

(G_d) 1, 11, 3, 4, 5, 10, 7, 8, 9, 9, 10, 12, 12, …

looks random, but who knows—maybe one can come up with a way of describing it that does the job. This is particularly pressing in view of the fact that there are many ways of describing a sequence: as we have seen, the description need not be num-

[8] If the alphabet contains $k > 1$ symbols, the number of all sequences consisting of *n* symbols from the alphabet is k^n, which is greater than $k^1 + k^2 + \ldots + k^{n-1}$ for all $n \geq 1$. Of course, if we only allow for certain sequences, i.e., if the sequences are coded in a language (say, the numerals) of lesser expressive power than the language available to describe them (say, English plus the language of number theory), then the result may not hold.

[9] See e.g. G. J. Chaitin, *Information, Randomness and Incompleteness*, Singapore, World Scientific, 1987.

ber-theoretic, and it need not be in the format that comes to us most naturally. That is the lesson of the Conway sequence. (In fact, coming to think of it, even (G_d) may very well be the beginning of a non-random sequence. The rule

(G^*_d) $S_n = n$ if the English name for this integer has an even number of vowels, otherwise $S_n = (n \bmod 5) + 9$

fits the data perfectly well…)

Now, why is this a problem? After all, Y knows that X *is not* thinking of a random sequence, at least if X is playing in accordance with R2, so Y knows that the sequence she has to guess *does* admit of a suitable description. Well, the problem is that Y cannot do much of this piece of knowledge. She knows that X is not thinking of a random sequence, but she doesn't know what this amounts to. Y doesn't know what sequences are ruled out because she doesn't have any effective procedure for telling what are the good candidates. She might believe that a certain way of continuing the initial segment is out because she might not be able to bring it under a rule—but she might just be wrong. Y might just be incapable of seeing the underlying pattern, so she might treat that as a random sequence when in fact it isn't. It might be precisely the sequence X is thinking about.

So here is the picture in a nutshell. If we stick to a generic notion of randomness as lawlessness, no sequence is random and R2 is perfectly useless. If, on the other hand, we stick to the more austere definition of randomness as incompressibility, then there are random sequences and R2 is fine as a matter of principle. Yet it is still useless when it comes to matters of practice. We may fully convince ourselves *that* a given sequence complies with the no-randomness requirement. But we may not be in a position to determine with certainty *whether* a given sequence complies with it.

5. Playing the Game for Real

So what are we left with? The "game of the rule"—it turns out—cannot be safely played. But why should that be a worry? After all, there are many other games that we can play, so why bother?

Well, there are indeed many other games we can safely play, but as we have already pointed out, this one is a game that we cannot dismiss so easily. The game of the rule is a game that *we play all the time*, whenever we engage in Humean inductive reasoning. We play it whenever we try to figure out the mechanisms of the world we live in, the laws of nature and the laws of society. We don't play it with number sequences but with the sequences of events that make up our history, when we try to make sense of them and see where they are leading to. And to say that this game cannot be played safely is not to say that we can stop playing it game cannot be played safely is not to say that we can stop playing it altogether.

At this point we have come to full circle and our story becomes a familiar one in the philosophy of science. Some think that the sort of skepticism that we have illustrated must be taken very seriously. For all we know, the world might not even be *trying* to play the game in accordance with R1 and R2, in which case our Humean inductive practices would just rest on a false presupposition. But even if the world were trying to play by the rules—even if the events with which we have to deal were fully in agreement with the presupposition of law-likeness—the fact that randomness phenomena cannot be effectively identified would be enough to justify a merely pragmatic attitude towards the endeavors of science. There is no way we can hope to "break the code". We can only hope to play the game in such a way that *we* find satisfaction in the laws that we envision, just as we find satisfaction in the social and political laws with which we try to regiment our daily interactions with our peers. Others feel differently. Not only do they think that we should play the game on the assumption that the world is issuing its challenges in compliance with R1 and R2; they also believe that we should not give up our hopes to get things right. After all, when we play the game for fun, we often win. We often succeed in uncovering the hidden pattern in spite of the difficulties that have been mentioned. Even if the initial segment that we are given does not uniquely identify the intended sequence, and even if we are not in a position to keep randomness under control, we often hit the correct rule because the other options are just too far-fetched to deserve serious consideration. So why not suppose that the same can happen when we play the game with the world of nature? All we have to do is to make sure that the world and we are on the same wavelengths, so to say—that what looks natural or far-fetched to us is indeed natural or far-fetched *simpliciter*. The history of science shows that sometimes we make mistakes, but that has never blocked scientists from pursuing their research with increased determination. On the contrary, the general thought has been that we can learn from such mistakes, and that we are getting closer and closer to winning the game at the next try.

This is no place to dwell upon this controversial dialectics. We may choose our party as we see fit. But the underlying predicament is something that can be best appreciated once we begin to see this familiar dialectics from the standpoint we have been suggesting here. For, on the one hand, we should not overlook one important sense in which playing the game with the world of nature may be easier than playing it *inter nos*. When both X and Y are people like us, each will try to win the game; in particular, X will try to issue his challenge in such a way as to make it difficult for Y to come up with the right guess—he will make every effort to design a rule that would be very hard, if not impossible, to discover in real time. By contrast, there is no reason to suppose that this is how the world out there issues its challenges. The world is not an intentional agent and does not care about "beating" us in the game, or so we may assume. In this sense, the practice of scientific induction need not be as hard as playing the game of the rule against a clever opponent trying to be smarter than we are, and the thought that we should try to be on the same wavelengths as the world is all but unreasonable. (That was indeed the main rationale behind Keynes's principle of "limited variety", a principle whose roots can be traced back to the philosophy of Francis Bacon:[10] an object of inductive inference should not be infinitely complex, nor determined by an infinite number of generators, and if we are assuming that the world is playing by the rules, we may well suppose that it is playing according to this additional principle, too.) On the other hand, there is also an important sense in which playing the game with the world of nature is not as easy as playing it *inter nos*. When we play, one player gives the beginning of the sequence and the other must figure out the rest. That may be tough, for the reasons that we have seen, but at least the input is clear—the first player is giving it explicitly. When we play with the world, by contrast, we must be careful. Not any series of events is on equal footing. We may witness the rising of the sun every morning and take that as an input for a law that we may reasonably formulate as a Humean inductive generalization. We might even think that it's worth looking at the series of events that we get by tossing a coin, for it might not be a random sequence after all. But when we zap channels during a commercial break, for example, the series of events that follow one another on our TV screen is not worth looking at. When we check the sky every time we hiccup, the series of events that we thereby collect is not worth any serious study. This is not to say that such series are random; there might even be a pattern, surprising as that might be. ("Every time I watch the Yankees they lose.") It's just that such series are not interesting. They don't count, so to speak. And they don't count because they would be there *even if it turned out that we live in a totally deterministic world*, a world where nothing is random and everything happens for a reason. To put it briefly, when we play with the world we have got to figure out *which* sequences to consider before we can start figuring out *what* they are, for the world does not tell us that. The world does not issue its challenges as explicitly as people do when they play with one another.

Now, precisely this is the main difficulty with Humean inductive reasoning. We can live with randomness and we can live with the fact that randomness is undecidable. Science has learned to cope with that, one way or the other. We can even assume that the world—unlike clever human players—has no interest in *beating* us. But we must be careful because our cognitive make-up is such that we constantly

[10] See J. M. Keynes, *A Treatise on Probability*; New York: Macmillan, 1921, Ch. 22. Compare the beginning of Bacon's *Magna Instauratio* and Book II of his *New Organon*.

look for patterns and trends even where there may be none. And there may be none, not because we may be dealing with random sequences, but because maybe there is no sequence to deal with. Maybe we are just exercising “the ridiculous art of Geomancy”, as Leibniz put it, and that’s not a way to play the game. For in the end, when it comes to playing the game for real, the one tacit Rule that we can *never* rely on is also the most obvious and important of all:

R3 *The initial segment of the sequence must be given explicitly.*