
Predictability and uniformity

This chapter expands on the issues raised in the previous chapter, and it tries to meet some objections that may be brought against the theses introduced there.

In the days and weeks following his first seminar (see Chapter 10), several of the Martian's audience dropped by his office to discuss some of the points he had made. Usually these exchanges were pleasant affairs, although a few of them did degenerate into heated arguments. In any case, it was obvious that interest continued – perhaps even increased somewhat – and someone asked the Martian to give some further seminars. This time he was not nearly so reluctant, even relishing the prospect a bit. And so, a little more than one month later, he stood before his colleagues again.

THE MARTIAN'S SUBSEQUENT SEMINARS

Several of you have talked with me about the theory I was trying to advance in my seminar six weeks ago. During these conversations, I jotted down some notes, and, in reviewing them for these additional seminars, it seemed to me that your worries and interests fall principally into two areas. Let me mention them at the outset, and then I will attempt to deal with them, one today, the other tomorrow.

The first area over which many of you expressed concern involves predictability. Several of you, certainly correctly, point out that much of what we choose to do is perfectly predictable, more specifically, capable of forecast on the basis of physical laws (or scientific law – the distinction is unimportant for this particular case), and you take this to show that physical laws can themselves be known temporally prior to their instances. But if this is so, then there are good

grounds for arguing that the truth-conditions of physical laws cannot (solely) reside in the events they subsume. In short, you allege that the empirical evidence of predictability confirms the Autonomy, rather than the Regularity, theory of physical laws. I will devote today's seminar to rebutting this argument.

The second point some of you urged on me was that the Principle of Determinism entails the Principle of the Uniformity of Nature, and that the latter – clearly – is inconsistent with our having done otherwise than we did. I will try to argue tomorrow that freedom to choose is not precluded by there being uniformities in Nature.

SEMINAR OF APRIL 12: PREDICTABILITY

My argument of last time was premised on the claim that the physical laws of human (and Martian) behavior take their truth from what we have done, are doing, and will do, rather than it being the other way around, that is, that we have done as we have, are doing as we do, and will do as we will do because the laws are what they are. Many of you have expressed your disagreement with this point of view, arguing that our behavior is in principle predictable and that this is so attests to the direction of logical priority¹ as running 'from' laws 'to' events:

You do not really choose at all. Still less does your choice determine the physical laws of the world. Your so-called choice is an action wholly determined by certain antecedent conditions (your genetic makeup, the energy transmissions your body has received over the course of your lifetime, etc.) and the physical laws of this world. If we knew everything there was to know about you, the information you have received and assimilated, the minutest details of your physiology, etc., and if we knew a sufficiency of physical laws, then your behavior, your so-called choice – however opaque your motives may seem to us in our present state of ignorance – would be perfectly predictable, in theory at least, if not in actual practice.

To begin, there is a certain ambiguity here in the word "predictable." If "predictable" is taken to mean "deducible," then my action – let's leave aside for a moment the question of whether that action might appropriately be deemed a choice – *is* predictable. After all, any event, E' -at- t_n ,

¹ [See "Epilog to the Martian's first seminar" in Chapter 10. –N.S.]

can be predicted (i.e., deduced) from two propositions of the form, “Whenever E -at- t_x , then E' -at- t_{x+1} ” (a physical law), and “ E -at- t_{n-1} ” (a statement of antecedent conditions). One can deny that events are predictable in this *logical* sense only at the unacceptable expense of denying the validity of *modus ponens*.

But this logical version of the claim about predictability is neutral on the question as to the order of logical priority. It is only on the temporal reading of “predictable,” the Laplacean epistemic version, that the claim may be brought to bear on the matter at hand. In this interpretation, “predictable” is taken to refer, as it more commonly does, to being able to say before an event occurs that it is going to occur.

What conditions must be satisfied for someone to make a successful prediction? A prediction of a future event will be deemed successful if and only if it is *true*. Thus a prediction will be deemed successful if it is true but irrational; if it is true but only a guess; if it is true and made contrary to the best evidence; if it is true but the result of a toss of dice.

We are not interested in predictions per se; our interest is directed to predictions that come about through the exercise of rationality, that is, to what may be called rational forecasts. All other foretellings will be disregarded. We exclude, too, whatever cases there may be of precognition and backward causation. Eventually, if their credentials can be established, we may wish to include these among cases of forecasts. I exclude them, not through any doctrinaire belief that they do not exist, but solely because they are, as yet, unproven to exist. And certainly, we will ignore prophecy.

What conditions must be satisfied for a successful rational forecast? Obviously, the prediction must be true. But more is needed. The prediction must issue from a rational inference either from known physical (or scientific²) laws and initial conditions or from ‘subliminal knowledge’ and observations. Admittedly, the latter characterization is vague. But it will do for our purposes. At any rate, I think it will be clear enough what I mean if I give an example. An experienced physician may be able to predict whether a patient will survive a medical crisis. Yet she does not do so by invoking any physical law or by painstakingly articulating the patient’s present condition. She does so by reacting to subtle clues in the patient’s appearance and

² As I mentioned in the introduction to today’s seminar, the distinction between physical laws and scientific laws is unimportant for present purposes. I will speak just of physical laws for the remainder of today’s seminar.

behavior and by having, although not necessarily being able to recount, an extensive firsthand knowledge of previous similar, relevant cases. Forecasts, then, may range from highly formalized arguments rehearsed internally (or even on paper, as in the case of weather forecasts), to the seemingly ‘intuitive’ (but misnamed) reactions of knowledgeable practitioners and experts. In any case, there are two components (besides truth) of a successful rational forecast: The forecast must ‘invoke’ universal propositions, and the forecast must call upon certain specifics of the matter at issue.

What follows from our ability to make rational forecasts? Does our ability to know of certain events, before they occur, that they will occur, imply that the universal propositions under which they fall have autonomous truth-conditions, that logical priority runs from physical laws to the events they subsume?

I want to argue that the Laplacean paradigm³ that claims that all events are in principle rationally forecastable is false. It is false if for no other reason than that not all of our behavior is forecastable, not even ‘in principle.’ And because it is false, it provides no selective confirmation of the Autonomy Theory.

There are two reasons why our behavior is not totally forecastable. One has to do with quantum indeterminacy; the other has to do with there being future events that are the first instances of their covering laws. I mention the first of these only to put it at arm’s length.

Events at the micro or atomic level sometimes do have causal results at the macro or phenomenological level. Witness Schrodinger’s cat (1983, p. 157) or the use of a Geiger counter to guide one’s looking for minerals hidden from sight beneath the ground. In these instances, the question of whether or not a certain large-scale, publicly observable event will occur depends, in turn, critically on whether or not some microevent triggers it. But if these latter microevents are genuinely unpredictable in the way in which current-day quantum mechanics suggests, then so too will be their macroscopic effects.

³ "We ought to regard the present state of the universe as the effect of its anterior state and as the cause of the one which is to follow. Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it – an intelligence sufficiently vast to submit these data to analysis – it would embrace in the same formula the movements of the greatest bodies of the universe and those of the lightest atom; for it, nothing would be uncertain and the future, as the past, would be present to its eyes." (Laplace 1951, p. 4)

For present purposes, this complication is a distraction and should be ignored.⁴ There are two reasons why this is so. For one, the fact of the matter is that the overwhelmingly greater part of our behavior is not infected by quantum indeterminacy. For example, a week ago the manager of my apartment building promised to repair the garage door when he returned from a trip on which he was then just about to embark. Yesterday he returned and fixed the door. The causal chain, which he had set in motion seven days earlier, lasted throughout the ensuing several days, and it would appear from the evidence that it was totally unaffected by the intrusion of quantum effects. Such sequences are common and familiar. Or consider how marvelously reliable and predictable are our computers. The probability of untoward behavior at the macro level (arising out of quantum disturbances at the micro level of their innards) has for all practical purposes been rendered zero by the sheer number of molecules making up their parts. Because of the nature of our own visual acuity, we tend to forget how truly massive, on an atomic scale, are many of the things that we cannot see with the naked eye and can observe only with optical or electron microscopes. A single nucleus of human brain cells may contain upward of 10^{12} atoms; or a transistor, invisible to unaided vision, residing on a computer microchip, may contain as many as 10^9 atoms. Granted, quantum events sometimes do make their indeterminacy felt at the macro level; but rarely in our central nervous systems or in our computers.

But even more important than the factual matter of the relative infrequency of the latter phenomenon is its logical irrelevance to our discussion. I want to argue against the Laplacean paradigm, but not because it makes the false assumption, which it does, that all physical laws are universal rather than statistical. I want to show that the Laplacean paradigm would still merit being discarded *even if there were no quantum indeterminacy whatever*.

For us to *know* beforehand, on the basis of a rational inference from a physical law, of any event whatever that it was going to occur, two conditions would have to be satisfied: (1) There would have to be a finite (and manageable) set of physical laws; and (2) each of these laws would have had to have already been instanced (so as to have been confirmed). Neither of these

⁴ [Stochastic processes and statistical laws form the subject matter of the next chapter. –N.S.]

conditions can be met. The number of physical laws is indefinitely large; and, daily, physical laws that have had no previous instance are instanced for the first time.

The completion of science has been a dream since the dawn of the age of Rationalism. It is a theme that ran through the seventeenth and eighteenth centuries. It motivated the editors and contributors to the *Encyclopedia of Unified Science* in the mid-twentieth century. And it reappears more strongly than ever as the driving force behind the search for a single theory to unify all the forces uncovered by physics. Nigel Calder, in chronicling the advances in physics in the twentieth century, titled his book *The Key to the Universe*. On its first page, he has written: “The key to the universe might be a *brief* set of equations or diagrams that encompassed all the large-scale and small-scale workings of the cosmos and showed their logical connections with the character of space and time” (1979, p. 7; emphasis added).⁵ And a few pages later: “Physics was always the master-science. The behavior of matter and energy, which was its theme, underlay all action in the world. In time astronomy, chemistry, geology and even biology became extensions of physics” (p. 14). The popular press, too, reinforces the belief in the deductive unity of science with their regular ‘news’ that a ‘breakthrough’ is imminent in physics that will complete Science. And physicists, like medical doctors, are often heard to remark that theirs is a profession dedicated to making themselves unnecessary: that the problems they address admit of eventual solution; indeed, that that solution is so close of late as to be virtually palpable.

This Laplacean idea that science is completable is a mistake. Inasmuch as some future human (and Martian) behavior will be the first instance of its kind or, to be more specific, the first instance of its covering physical laws, then those covering laws could not, could not even in principle, now be confirmed. Hence they could not now be known, prior to the occurrence of these future events.

No one can deny that much of our behavior *is* capable of being forecast, particularly our ‘routine’ behavior. That which we do regularly lends itself especially well to being forecast. Anyone who knows me knows that I regularly have a cup of coffee at mid-morning. You reliably forecast that I will tomorrow. Yet that I will hardly threatens my freedom to desist if I should

⁵ [Recall the quotations from Carnap and Schlegel cited in Chapter 1, footnote 9. –N.S.]

want to. If I wanted to frustrate your prediction, I could; but I do not; indeed, I want that it be true. So you predict that tomorrow I will freely choose to drink a cup of coffee, and tomorrow I freely choose to drink the coffee.

That some event lends itself to being forecast and also is an instance of someone's exercise of her freedom is perfectly consistent. But this does not mean that *every* exercise of one's freedom is in principle capable of being forecast.

There are (at least⁶) two kinds of case of free action that are incapable of being forecast. Michael Scriven (1965) has written about one of them in detail, namely, the case of the contrapredictive, the person who sets out to frustrate the forecasts made about his behavior (my example: Eisenhower planning the invasion of Normandy). I will not pursue this case.

The other kind of case is the one mentioned in the first seminar, namely, the case of a person's being in a unique situation, for example, being the one and only person ever to choose which mortality table (merged-gender or sex-distinct) society is to adopt henceforth.

Of course every situation is, if given a complete enough description, novel; likewise, every situation is, if given a lean enough description, routine. When we come to make forecasts, we can never give 'complete' descriptions unless we describe the entire universe at that moment. Realistically and practically, we make do with 'partial' descriptions, and for many (most) forecasts these partial descriptions suffice. For example, the light switch's being flicked up is followed by the chandelier's lighting. And it matters not one whit whether the lever on the light switch is white, brown, red, or blue. Something less, indeed something *very much* less, than a complete description suffices for a successful forecast in this case.

But electrical circuits consisting of a source of electromotive potential, a light switch, and a chandelier are not entities endowed with free will. Forecasts about entities endowed with free will are more risky, not just because the number of variables and physical laws are greater, but also because some of the laws under which the behavior of the entity may be subsumed are a matter of that entity's choosing.⁷

⁶ There may well be more. No other obvious cases spring to mind, however.

⁷ Typically, our formally crafted forecasts utilize several laws. It will simplify matters, however, if we conceive of forecasts as issuing from a single law. There is no harm in making this simplifying assumption, for any set **S** of universal propositions that, in conjunction with a set of statements of antecedent conditions, logically implies a forecast *F*, will also logically imply the existence of a single universal proposition L_i that together with the set of antecedent conditions similarly implies *F*.

Suppose that up to today I have always bought brown shoes instead of shoes of any other color. If I continue to buy brown shoes for the rest of my life, then the following proposition is a physical law: “Whenever a person having the properties ... (here one gives a unique description of me in purely descriptive terms) ... buys shoes, he buys brown shoes.” Further, suppose tomorrow I need a new pair of shoes and go to a local shoe store to purchase a pair. At that moment, it is up to me to choose (although this probably will not be the intent of my choosing) how novel – as regards forecasting my choice – that situation is to be. If I choose black shoes, that situation will be more novel than it would have been had I chosen brown; the novelty of this situation will vary according to what I in fact *do*.

We look at the past behavior of things and try to formulate, using the fewest number of predicates we can get away with, universal propositions that are found empirically to provide grounds for forecast.⁸ Many times, this works admirably (e.g., the case of the chandelier, the case of my drinking coffee tomorrow morning). Other cases are riskier: You predict, on the basis of my previous exceptionless behavior in this regard, that I will buy brown shoes on the next occasion of my purchasing shoes. You may well be right. If I *do* purchase brown shoes and happen to die immediately thereafter without ever having bought another pair of shoes, the universal proposition on which you based your inference would turn out to have been a physical law. If, however, I should happen to frustrate your forecast (either deliberately or unknowingly) by purchasing black shoes, that same universal proposition would turn out not to be a physical law. Either way, your forecast would have been rational. But whether you base your prediction on a physical law or on a false universal proposition lies in *my* hands. It is up to me to make that universal proposition in your inference a physical law or no law at all.

Forecasting what color shoes I will next purchase can call, at least, on some previous behavior of mine (and of some other persons as well) regarding shoe purchases. But my choosing for society between two kinds of mortality tables is without any precedent at all. Short of describing it as “a case of choosing mortality tables for society” (which description has no prior

⁸ I am not suggesting that our laws are Baconian generalizations from experience, although of course some few do originate in that way. My point is one about justifying universal propositions, however we may come to entertain them in the first instance.

instance whatever), there does not seem to be any broader (less specific) description that carries with it much, if any, clue as to what I will do. Obviously, we can describe it as “a choice,” or even as a “monumental choice,” or as a “choice in which principles of fairness are seen to be at play,” etc. But none of these plausibly provides any clue as to what I am going to decide. Suppose that once before in my life I was called upon to make a monumental decision. Suppose I was asked to choose what the punishment was to be for someone’s killing all the surviving members of an endangered species. Thus there would be a physical law to the effect: “Whenever ... (uniquely individuating description of me) ... is faced with a monumental decision, *JKL*, he chooses *PQR*.” If we narrow the description *JKL* so that it applies just to the case of my choosing the fit punishment in the ecological case, then we can give a fairly specific, narrow interpretation to *PQR*: namely, “that the culprit should be strung up by his/her thumbs in a public square.” But if we use these interpretations of *JKL* and *PQR* respectively, then the physical law just simply does not apply to the case of my choosing which mortality table is to be used. But if, on the other hand, we broaden *JKL* so as to be general enough to include both my monumental decisions, that is, the one about the earlier ecology case and the one about the upcoming insurance case, then we are at a dead loss as to what to make of *PQR*, for we have no independent way of knowing what my second monumental decision is going to be.

We can imagine the following kind of objection.

But couldn’t one give the event you are trying to predict a broader description such that that less specific description *does* have prior instances, and thus one would have inductive warrant for the appropriate conditional statement?

The answer: No. And the reason is that the description, “chooses *B*” (i.e., “chooses that merged-gender mortality tables be used”), already is the minimum useful description. We have already ‘peeled off’ all the other aspects of my upcoming decision that are of no particular interest: for example, whether, when I announce that decision, I am wearing a brown suit or a blue one; whether I speak in a firm, commanding voice or a shaky, nervous one; whether I glance at my friend in the press corps or avoid her eyes; etc. In short, there just is *no* broader description of what you are trying specifically to predict. And although there are, of course, less specific descriptions that will *fit* my action, for example, “will announce his choice,” these descriptions

are too broad for your purposes. You are not content to predict simply that I will choose *something* or even that I will choose between two specific alternatives. You want to know specifically what that choice will be; and that is just what it is that is without prior instance.

My point about forecasting events is really quite simple: Our ability to forecast depends on there having been prior positive instances of the relevant covering law.

Most of the physical laws that are conveniently lumped together as belonging to the domain of physics have already been instanced. Thus, without there ever having been a silvery blue meteor striking Earth, we can rationally forecast the impact of one when we see a silvery blue meteor through our telescopes and plot its path as being on a collision course with Earth. In this case, the relevant laws of orbital motion are known to us, and experience has taught us that the motions of celestial wanderers are independent of their colors.

Then, too, much of our routine behavior lends itself to rational forecast. That is why we call it "routine." There is an analytic connection between "being routine" and "being forecastable."

But these two bodies of forecastable events should not lead us to believe that all events are in principle capable of being rationally forecast. Acts of choosing – my habitual coffee drinking notwithstanding – provide countless examples of events incapable of being rationally forecast.

I am not here invoking the fallibility of inductive inference; although that is of course true, it is not my point. Rather, I am suggesting that there is nothing known prior to my choosing that merged-gender mortality tables are to be used that provides any grounds at all for rationally predicting that that will be my choice. After the fact, any number of universally true propositions describing my choice will be knowable. But none before. I am thus hardly making a skeptical point about the possibility of knowing something before it happens. On the contrary, I have already allowed that all sorts of events are knowable in advance of their occurring, even though our belief that they will occur follows in part from universal propositions not all of whose instances have occurred, that is, from universal propositions at least one of whose instances is the forecast event itself. Indeed, it should be clear that one can know some future contingents even when those propositions are not epistemically certain; even when it is logically possible that they should be false; even when it is logically possible that the universal proposition(s) figuring in their forecast 'may be false'; etc.

Those of you who subscribe to the Laplacean paradigm may remain unconvinced. You may try to argue this way:

The predictor wants to be able rationally to infer from knowable data the proposition “Arbiter *a* will choose merged-gender mortality tables.” Why should he be unable to do this? What makes this case so different from other cases of forecast? Although the predicate “chooses for society merged-gender mortality tables” is without previous instance, the two predicates of which this is compounded, namely, “chooses for society” and “chooses merged-gender mortality tables,” have each, after all, independently applied to former cases of choosing. Legislators, jurists, arbitrators, et al., have often made choices for society. And managers of various insurance companies and pension plans have chosen that their firms and plans would use merged-gender, rather than (as do other firms and plans), sex-distinct tables. Is it really impossible to find in the histories of these respective classes of decisions a pattern of behavior that would allow one rationally to predict what will be the arbiter’s decision in his choosing which of the two will henceforth become the single, mandated standard?

I think it is impossible to *find* the pattern. The trouble is that there are too many data available for too few cases. Suppose there have been six occasions when the managers of insurance companies have chosen that their firms would use the merged-gender tables, and an equal number of occasions when managers of other firms have chosen the sex-distinct tables. (The exact number is of no especial moment; neither is the fact that there are equal numbers on both sides. I choose the number six only because it is convenient.) What universal propositions are we in a position to confirm? Simply, a very great number. Given the extraordinary complexity of managerial decisions of this sort – the number of factors considered by the choosers, their physiologies, their temperaments, their manners of consulting one another, their biases, their histories, their information, their ... (continue ad libitum) – there will be an indefinitely large number of confirmable universal, conditional propositions, all of which have as their consequents “chooses merged-gender mortality tables” but will differ in their antecedents. Some of these will have relatively short antecedents, that is, will be relatively broad, and others will have lengthy antecedents, that is, will be relatively specific. At the very least, one (provided determinism is true), and far more likely, a very great number, will apply to the case of my

upcoming decision, and were they to be identifiable as *true* could be used successfully to predict what I will choose. But can these true propositions be identified among the myriads that, although consistent with present data, will be proved false by future events? Consistency with present data is not the same as pervasiveness, and it is evidence of pervasiveness that is needed for a rational forecast.

Suppose the six managers who chose the merged-gender tables were all college educated and more than fifty years of age; and suppose the six who chose the sex-distinct tables were all owners of sports cars and members of tennis teams. Now I happen to be less than fifty and own a sports car and am a member of a tennis team. Shall we predict that I will choose the sex-distinct tables? Before we do, we should note that the six managers who chose the merged-gender tables all are classical music devotees, whereas the six who chose the sex-distinct tables are jazz aficionados. I prefer classical music. Shall we then predict that I will choose the merged-gender tables? The universal generalizations that are compatible with our data pull us in opposite directions. We cannot rationally choose between these competing generalizations. We know that once I have made my choice some of these generalizations that are consistent with present data will not be consistent with the enlarged set of data that includes the statement as to what I choose. That is to say, some of these competing generalizations that are consistent with present data are not timelessly true; future events render them false. But, from where we stand, here and now, prior to those events, we cannot tell which of these competing hypotheses, leading to contrary predictions, are timelessly true and which are timelessly false.

It always has been thus. Francis Bacon argued that perhaps God could discern the truth from a collection of positive instances, but mortals need the winnowing device of counterexamples. As the number of variables, or factors, or predicates needed for successful prediction increase, then so too must the number of cases so that we should reasonably be able to select among competing generalizations. We already know enough about this world to know that we cannot predict human behavior on the basis of a few factors, for example, on place of birth, color of hair, weight, blood type, number of books read, etc. Human (and Martian) behavior is just too complex. Predicting it requires considering a great number of factors. But as the number of factors increase, the number of their instances decrease, often so much so that it is impossible to get a fix on which of the competing generalizations is true. Surely there must be such cases. Whether or not my example of choosing mortality tables finally is such a case is less important

than that there be *some* such cases. And it is reasonable to believe that there are. There seem to be so many factors to be considered in predicting human behavior that to believe that every kind of behavior has already been instanced flies in the face of the empirical facts.

There is a final point to consider. It may occur to someone that the argument about the unpredictability of my choice may rest upon trying to predict what I will *say*, and predicting verbal behavior is notoriously difficult. One may be inclined to argue this way:

Suppose the legislature that entrusted you with the decision furnished you with two postcards, already addressed and filled out, one for each of the two possible decisions. Your decision, as to which of the two mortality tables society is to use will be made by your picking up the one postcard on the left or the other on the right. Your action will be constituted by the physical motion of your moving your arm to the left or to the right. And surely *that* is predictable. If one describes what you are doing as 'choosing' or as a conscious action, then under that description the event may well be unpredictable; but inasmuch as what you do is also a physical movement, it must then be predictable just as are other physical movements, for example, the operation of a pop-up toaster or the turning on of a clock radio.

A retreat to mechanism or physicalism will not help. It just begs the question. For clearly I will want to argue that the movement of my arm to the left or to the right is forecastable only through the use of a universal generalization whose antecedent is a complex clause, so complex in fact as to render that generalization, like the previous ones pertaining to human action, also without inductive support. My argument for unpredictability turns on the complexity of this world and, as a consequence, on there being competing generalizations which, although consistent with available data, are inconsistent with future data. My argument does not turn on any postulated or supposed metaphysical difference whatever between conscious and unconscious behavior, nor on any difference between human (or Martian) *action* on the one hand and mechanical or physical or autonomic behavior on the other.

If all events were in principle forecastable, then one would have powerful grounds for believing that physical laws bear their truth-values autonomously. But the belief in universal forecastability is not warranted by empirical data; and if it is fostered by the Autonomy Theory itself, it can not then be used in turn to bolster that theory.

Available contemporary empirical data is a two-edged sword in this debate. Physics and chemistry have had stunning successes in predicting behavior for certain kinds of events. And generalizing from this success in the natural sciences, many persons have come to think that universal forecastability applies, in principle, within behavioral science as well, and to the case of human behavior in particular. Failures in predicting human behavior are – according to persons attracted to the Laplacean paradigm – to be explained away by our lack of imagination in creating new hypotheses, to inattentiveness to details, to failures of perception, to errors of inference, etc.; in short, in one way or another to our not selectively confirming our hypotheses by finding what is there to be found.

But the failure to predict human behavior need not be looked upon as evidence of our shortcomings in doing empirical research. For it could be that the fault lies not in us, but in Nature: that there is nothing there to be known prior to the occurrence of the event we would like to predict. As one's science passes up the ladder of evolution – looking first at atoms, then at molecules, at proteins, at protozoa, at ... (eventually) human beings (and Martians) – the laws under which behavior is subsumed become more and more complex and more and more numerous. And, thus, surely it is reasonable to assume that many of these timelessly true laws covering the latter products of evolution have no instances as yet, that their first instances will occur a minute, an hour, a year, or a millennium from now.



In summary, I am under no illusions that I have disproved Laplaceanism. I do not think it can be disproved. Or, to be more exact, those persons who do not share certain other of my beliefs are bound to regard my objections to Laplaceanism as question begging. These background beliefs are, however, very much in contention, most principally the claim that we are sometimes free to choose among genuinely alternative courses of action.

What I am doing is not so much trying to disprove Laplaceanism as to provide a viable alternative to it. Determinism – the theory that every event is a second member of a sequence that falls under a physical law – is, so far as we can tell, a reasonable hypothesis to have about this world, at least at the macroscopic level of description. But Determinism, thus explicated, is compatible with two theories: the Autonomy Theory, which gives logical priority to physical laws and sees events unfolding so as to accord with these laws; and the Regularity Theory, which assigns logical priority to events and sees physical laws as taking their truth-values from whatever happens in the world.

The Autonomy Theory is a virtual consequence of Laplaceanism. If one believes that in principle the future course of the world (again disregarding quantum effects) is forecastable, then one must believe that there really is a finite, manageable set of physical laws that when taken in conjunction with a statement describing the physical universe at any one moment of time (or on two successive moments⁹) does imply the entire future course of the universe. But even if we disregard quantum indeterminacy, there is no good reason to think that this universe is at all like this.

Unlimited forecastability of the Laplacean sort can occur only if what happens in the future is not truly ‘open’ as to possibilities. If everything that happens must accord with physical laws, then the fewer the number of physical laws, the fewer the possibilities open to us. For our choices to be genuine – not idle shams of ‘going through the motions’ and then finally opting for what was inevitable in any case – there must be many possibilities open to us, and this in turn is possible only if there are many physical laws. Freedom of action is compatible with Determinism; but it is not compatible with Laplaceanism because Laplaceanism is premised on the belief that there is a finite, manageable number of knowable physical laws. If we are free on occasion to choose among genuinely alternative courses of action, and if Determinism is true, then (1) there is no fixed upper bound on the number of physical laws, and (2) the number of instanced physical laws, and hence knowable physical laws, will steadily increase with time. If Determinism is true, and if we are free, then Laplaceanism, that is, forecastability in principle, is false.

SEMINAR OF APRIL 13: UNIFORMITY

At the time it was first mentioned,¹⁰ I said that perhaps the most paradoxical-sounding corollary of the Regularity Theory is this:

If a person is in a unique situation and chooses *B*, then his choosing something else that is logically inconsistent with *B*, for example, *C*, is physically impossible; nonetheless, his choosing *C* may well have been unconditionally possible.

⁹ The qualification about two successive time-slices is required because of the existence of laws whose antecedents refer to the histories of situations rather than to momentary states, e.g., Steinmetz’s Law of Hysteresis. (See Huntington and MacCrone 1983.)

¹⁰ [See “The Martian’s first seminar” on page 127. –N.S.]

The way I suggested to remove the seeming air of contradiction is to remember “that according to the Regularity account, ‘physically impossible’ means nothing more or less than that such an event (e.g., choosing *C*) never, ever occurs. But ex hypothesi, this is true; for remember as well that we have ... stipulated that the situation in which the chooser opts for *B* is *unique*. It is logically impossible that he should choose both *B* and *C*; so whatever he chooses, the other would thereby become physically impossible. Still, he could choose – unconditionally – the physically impossible action if he wanted to. It is his choosing one of the two alternatives that makes the other one physically impossible. But he could choose either.”

Now in our subsequent private discussions, some of you have returned to this issue and have argued that one’s choosing *C* is not possible, that if a choice is genuinely determined then no other choice can occur, and that this is guaranteed by the Principle of the Uniformity of Nature. You have argued that if and only if Determinism holds, then the Principle of the Uniformity of Nature holds; they are one and the same principle.

Laplace had stated the Principle of the Uniformity of Nature in this way (although he did not give it any name whatever):

If the present state of the universe was exactly similar to the anterior state which has produced it, it would give birth in its turn to a similar state; the succession of these states would then be eternal. (1951, p. 171)

And Mill expressed it this way:

There are such things in Nature as parallel cases, that which happens once, will, under a sufficient degree of similarity of circumstances, happen again. (1965, p. 201)

There is one interpretation of this latter dictum that is relatively harmless. It is the one that would make of it a methodological principle counseling that one never abandon the search for similarities, in effect a truism invulnerable to refutation:

If seemingly ‘similar’ events fail to have identical outcomes, then the similarity was not in fact extensive enough; there was some relevant difference that we have overlooked, and further empirical investigation will reveal it.

It is, however, another interpretation of these dicta that occasions the present inquiry.

Clearly, what is common in both Laplace's and Mill's versions of the Principle of Uniformity is the notion: "Same history, same future." It is implicit in Mill's version and explicit in Laplace's that, were the world to be 'backed up', it would have to have the same future as the one it in fact has had, that there was but one possible future for the world, given its initial state. Applying this notion to the present case, one might argue:

Whatever led up to one's choosing *B* could not have been followed by *C*. Were the world to be brought back to its state just prior to one's choosing *B*, one would choose *B* again (and again, and again, for however many times the world were to be backed up). In this world, the state prior to *B* could only be followed by *B*, never by *C*.

We cannot, of course, conduct the experiment. We can't back up or reverse this world and approach a decision point for a second time to see whether the same outcome will occur the second time through. The best we can do is to postulate another possible world identical in its history to this world up to some time *t* and ask whether the Principle of Determinism really requires that that world have the same future, subsequent to *t*, that the actual world has. I claim that the Principle of Determinism requires no such conclusion.

The maximal amount of similarity possible, for the case under consideration, is another world, *W'*; identical in its history to the actual world *W_a* up to the moment *t* but in which at *t* I choose *C*. Both are worlds in which the Principle of Determinism is true: Their respective events are members of sequences that fall under physical laws. World *W'*; *could have been* our world, if only I had chosen *C* at *t* rather than *B*. In short, the Principle of Determinism does not entail "same history, same future," does not entail that a given history can have one and only one possible outcome. A given history cannot, of course, have more than one outcome; but which outcome it will have is not something predetermined by that history. True enough, the future of that world is logically implied by the description of its history together with that world's physical laws. But, to repeat, there is no predestination in any of this provided those laws take their truth from what has happened, is happening, and is going to happen, and not from some other source. Provided the laws are not autonomously true, that is, are not logically prior to their instances, there is no predestination: Nothing 'compels' or 'requires' events to occur. In a categorical sense, within empirically discovered limits, we are free to choose from among a set of genuinely

alternative actions. How *much* we can choose, and what is outside our sphere of choosing, is something we *discover* about this world; it is a contingent matter knowable only a posteriori. But the crucial point is that there are some, at least, genuine choices open to us.

The Principle of Determinism and the Principle of the Uniformity of Nature are *different* principles; the latter is a stronger principle in that it implies the former. The Principle of Determinism does not imply the Principle of the Uniformity of Nature (no more so than that it implies forecastability in principle). (Note that the Principle of the Uniformity of Nature is often stated counterfactually. Recall, for example, Laplace's statement quoted a moment ago. The Principle of Determinism states merely that events fall under universal laws. It does not state that events can have only one possible outcome.)

A further difference between the Principle of Determinism and the Principles of Uniformity lies in the fact that *being determined* and *being uniform* are different concepts. This can be seen by noting that the two fall under quite different determinables. *Being determined* allows of only two different values: Either a situation (event, etc.) is determined, or it is not. But *being uniform* admits of degrees: One situation may be more, or less, uniform than another.

Many different analyses of "uniformity" can be given. But one is appropriate for present purposes, and it is a comparative concept; not a classificatory or a quantitative one¹¹:

A world W_1 will be said to be more uniform with respect to the occurrence of an event-kind e than another world W_2 if a set of minimum conditions sufficient in W_1 for the occurrence of an event of kind e constitutes a proper subset of a set of minimum conditions sufficient in W_2 for the occurrence of an event of kind e .

Intuitively, this means that it is easier (ontologically speaking) to bring about an event of kind e in W_1 than in W_2 : Fewer conditions must be satisfied in W_1 to bring about events of kind e . But note, W_1 and W_2 may both be perfectly deterministic worlds.

Allow me a few illustrations. Consider, first, what happens when one plays a flame on the surfaces of various objects in this world. Some objects will ignite; others not. Whether they do or not depends on a variety of factors, including their chemical compositions, their abilities to conduct heat rapidly away from the point of application, etc. But in some possible worlds, fewer

¹¹ For more on these three types of concept, see Carnap 1962, pp. xv-xix, 8-15.

conditions for igniting must be satisfied than in this world. We can imagine a world, for example, in which the thermal conductivity of the object has no role in whether or not it ignites. In that world, the ignition of objects would be more uniform than in this. Similarly, we can describe a possible world in which all metals have the same electrode potential. Such a world is more uniform, as regards electrode potential, than a universe (such as this one) in which the electrode potential varies from metal to metal. A universe in which all round things are hot and all square things are cold is more uniform, as regards the relation between shape and temperature, than a universe (such as this one) in which some but not all round things are hot and some but not all square things are cold.¹² And of course, on the other side of the coin, a universe in which the various isotopes (atomic weights 192, 194, 195, 196, and 198) of platinum (atomic number 78) are yellow, blue, green, blue, and red, respectively, is *less* uniform (as regards the color of platinum) than a universe (such as ours) in which all isotopes of platinum are silver colored.

Obviously, just about all (perhaps all) connected event-kinds in this world are both less uniform than they logically could have been and – equally – more uniform than they could have been. How *much* uniformity there is in a world,¹³ and how it distributes itself over the class of connected event-kinds, will vary from possible world to possible world.

Certainly, there are uniformities in this world. But the crucial question concerns the degree of uniformity and whether the uniformity is of such a degree as to threaten free will. It is at this conceptual juncture that the question has become confounded by confusing it with the question about physical determinism.

One of the main motives in the rebellion against determinism, not only on the part of ordinary people but also of modern philosophers who have been most vigorously opposed to the determinist position ... [is] ... the desire for freshness, novelty, genuine creation – in short, an open rather than a closed universe. ... even a good year – had we to live it over again with every detail fixed beforehand – would stifle us with boredom: our food would taste dull to our palate, our most spontaneous talk sound as uninspired as the playback of a tape-recorded conversation, and our words of love would sound hollow because we should know beforehand the precise moment of fatigue when they would expire. (Barrett 1958, p. 31)

¹² Indeed, this case comprises a limiting case of the explication being considered, since there seems to be no connection between shape and temperature in this world.

¹³ An unexplicated, intuitive concept of the ‘amount’ of uniformity is adequate for present concerns.

There is a curious inconsistency in Barrett's scenario: We are asked to imagine living a good year "over again with every detail fixed beforehand" and are told that were we to do so we would be "stifled with boredom." But unless we were stifled with boredom the first time around, we could not be stifled with boredom the second time around, not if – as he says – the second time around every detail is 'fixed' (i.e., a repeat of the first time around). For if we were not stifled the first time around, then, were we to repeat exactly *that* experience, logically we could not be stifled the next time around either. Obviously, Barrett has something much less than total or exact repetition in mind. He imagines, rather, repeating an experience but *knowing* it to be a repetition; and *that* might well be boring. I am not quite convinced, however, that it certainly will be boring. I remember hearing many years ago a performance of Bach's "Third Sonata for Unaccompanied Violin" that was utterly sublime. I would not mind in the least, indeed would welcome, a repeat of *that* experience. In any case, I think the thrust of Barrett's statement is real enough, and if it is not altogether a perfect piece of psychologizing, it is, nonetheless, probably not too far off the mark.

But what if experiences were repeated exactly, complete with all of one's feelings of anticipation, deliberation, expectancy, and concern? In such circumstances, could our actions be said to be free?

Imagine, if you will, a possible world – mercifully unlike this. one – where daily, throughout all eternity, every man at 7:32 P.M. closes his evening paper, folds it into thirds, clears his throat, and then challenges his wife (every man has exactly one wife to whom he is wedded forever) to a game of chess. Every day, he thinks carefully for a moment and then makes the same opening gambit, and his wife responds with one and only one defense, which soon changes to the offense, with her placing his king in checkmate in thirty-two moves completed at precisely 8:47 P.M. Gruesome world, isn't it? One only a Martian mathematician, those reputedly most regular of creatures, could relish.

My visceral feeling, my prephilosophical, untutored intuitions tell me that the unfortunate denizens of that repetitive world neither exercise free will nor even for that matter *have* free will, that any seeming 'deliberation' they go through prior to 'choosing' to play chess is illusory, a counterfeit of the real thing.

Can I do better than to call upon these emotional reactions, reactions that, possibly, some of you do not share? Could it be that there is really nothing more wrong with such a world than that I find the prospects of living in it odious? Is it really, as I allege, devoid of free will? Couldn't one challenge my claim by arguing:

There might, after all, be free will in that world. Perhaps all the benighted folk in it carefully, deliberately, choose to play chess. They did not, and do not, *have* to; but they *choose* to. By the free exercise of their wills, they choose, day after day, to lead lives of a repetitive sort. They could do otherwise, if they should want to; but they do not want to.

Surely, one thing I cannot say in rebuttal is that these chess players are driven by determinism, or must act out their dreary lives in accord with powerful physical laws of their world. I do not want to give such an account for *any* possible world, it being totally antithetical to my analysis of physical law. Rather, my rebuttal must take quite another form.

Let's approach this problem in a two-step process. Let's begin, not with the more difficult problem of whether there is, or might be, free will in such a repetitive world, but with the easier problem of whether there is any *evidence* of the exercise of free will.

We can imagine our pairs of chess players having the following conversation after their matches:

He: Why did you move your Queen's Knight to QB4?

She: I anticipated that your attention was directed to your setting up the Baron Scarpia Gambit, and I figured you would not notice that I was implementing Boskovshy's Ploy. It was taking a bit of a chance I know, but I counted on your being distracted, and I chose to go ahead with it.

He: Bad luck. If I had chosen instead the Stonewall Maneuver, as I was seriously considering, you never would have been able to get the upper hand.

From Olympian heights, eavesdropping on this possible world for the first time, we would probably be inclined to take this conversation at face value and consider that it provides evidence that the speakers are creatures endowed with free will. But when we notice that they have the identical conversation on the following day, and again the day after, and still again on the fourth day, indeed on every day, their words sound empty, as they would in our watching a play being acted for the umpteenth time. The more often the speakers recite these lines, the more they come to resemble Walt Disney's automatons (trademark: "animatronics").

Introspective feelings of freedom, public avowals of freedom, and the like are all defeasible evidence of true freedom. Other kinds of evidence can prompt us to discount such *prima facie* evidence. Compulsions, for example, may go unrecognized by their owners; drug induced behavior, such as the startling effects (megalomania) of cortisone, may be completely unrecognized by their exhibitors; etc. And so it is with the repetitive world we are considering. The avowals of freedom on the part of its inhabitants do not ring true. Indeed, the continual repeating of their avowals eventually is construed as evidence against the very thing being claimed: “Methinks the man doth protest too much.” Their claims come to take on the aspect of self-frustrating or self-refuting reports. Apart from their avowals, which we would be inclined to dismiss, nothing whatsoever suggests that anyone acts freely in that world. All the evidence of that world suggests that the persons therein do not exercise any free will.

We now come to the second stage of my rebuttal.

If we agree that there is no good evidence that the persons in the repetitive world are exercising their free will, couldn't someone nonetheless argue that they *might* be?

Nothing that happens in the repetitive world suggests that the persons in it are exercising any freedom. But then again nothing suggests that they are not. It is possible, it is certainly consistent with the evidence, that they *are* exercising their free will.

This defense is problematic on several counts. First of all, I think that total lack of evidence *for* something (e.g., the exercise of free will) ought to be taken as evidence that that thing does not exist. But this is, admittedly, merely a methodological principle to guide our apportioning rational belief. It hardly clinches the matter metaphysically speaking. But still, it should put us on our guard, for it suggests a second point.

Suppose one were to allow that this repetitive world really, contrary to the evidence, was one in which the inhabitants did exercise their free will. How – we now ask – would such a world differ from a world identical in all appearances but in which the inhabitants did *not* exercise their free will?

In other words, if the exercise of free will does not ultimately come down to exhibiting certain kinds of behavior, in what then does it consist? We can imagine someone who wants to defend the position that the repetitive chess players might really be exercising free will, saying:

The husband was exercising his free will because, as he said, he seriously considered the Stonewall Maneuver, but *chose* not to pursue it. He could have done it if he had wanted to; but he simply did not want to. In some other possible world, very like his own, he did invoke that maneuver, however.

This will not do, and the reason why lies in the fact that the alleged difference between the repetitive world in which a person does exercise his free will and a repetitive world in which that person does not exercise his free will is no distinction at all. That that person does something different in some *third* possible world, totally fails to distinguish between those two possible worlds in which the chess players act absolutely identically but for one difference: In one they are exercising their free will and in the other not. And if one tries to say that this third possible world, the one in which the chess player does implement the Stonewall Maneuver, is accessible from the former but not from the latter of the two otherwise indistinguishable worlds, we are at a loss to understand what difference this talk of accessibility is supposed to denote. Indeed, it is nothing other than the very difference one is intent to explicate masquerading in the elaborate costume of the possible-worlds idiom.

Beware. The question at issue is not whether there is some possible world ‘very like’ the repetitive world, but in which the husband’s actions do issue from the exercise of his free will; but rather, whether, under the specialized kinds of circumstances described (unrelieved repetition), his actions could be said to be free. The question is not, then, whether there is some possible world in which the man’s actions (described as simply “foregoing the Stonewall Maneuver”) are free – of course there *is* such a world; but whether there is a possible world in which *both* his actions are free *and* his behavior perpetually repeats. Although I would want to insist on a positive answer to the first question, I am equally inclined to answer in the negative to the second.

It would be natural for you to construe my remarks as evidence of a certain verificationist tendency on my part, and I am not altogether sure that at root they might not be. But if they do issue from that source, they do not do so in quite the straightforward way you might at first assume.

My inclination, as you have seen, is to say that the repetitive world I have described is one in which the inhabitants are not exercising free will. And you might assume that I reach this conclusion on the basis of there being no evidence for their exercising their free will. But, you might object, they really might be exercising their free will, even in the absence of there being any evidence:

Consider a possible world in which there are invisible elves who manage always to escape our detection. We would be inclined to say such elves do not exist. But they really do. Lack of evidence for a contingent existent is logically consistent with that thing's existence. Lack of evidence for a thing's existence can be taken as good reason to disbelieve that that thing exists; but it cannot be taken as proof of that thing's nonexistence.

As natural a rebuttal to my position as this objection is, it does not quite hit the mark, for there is one all-important disanalogy with the situation as I am imagining it. You see, in the world where there exist xenophobic elves who forever elude our observation, we can imagine someone or something else detecting their existence. Presumably, the elves themselves know they exist; so too does a god whose powers of observation surpass yours and mine. More exactly, what sets the elf-inhabited world apart from the repetitive world I earlier described is that there is something *more* to be *observed* in the elf-inhabited world. There is a 'detectable' fact of the matter, which somebody or other, not you or me, could learn. In the repetitive world, there is no such 'additional' fact. I take myself to have given you the entire relevant data base. The problem with the repetitive world lies not with our lacking a further fact, but with our interpreting the facts as we have them. As I am imagining it, not even a god of superior powers of observation would be any better off than we in deciding whether these chess-playing couples are exercising free will. The problem is a *conceptual*, not a scientific or epistemological, one. I cannot believe that what accounts for the difference between the exercise of free will in a 'spontaneous' game of chess and the lack of free will in a daily game of chess comes down to something we could hope to find in the situation itself, either in, for example, the brain-states of the actors or in the immediate externals.

It seems to me that, of two similarly appearing games of chess, one occurring in a world where chess games are pretty irregular, and one occurring in a world where chess games are as regular as the decimal expansion of $5/33$ (i.e., 0.1515151515 ...), the former game might

well and truly be deemed to issue from the exercise of one's free will, whereas the latter would never be so.¹⁴

The difference has to do with the very fact of repetition. Highly repetitive acts, even when accompanied by avowals that they issue from the free exercise of one's will, seem – just by the fact alone of their repetitiveness – not to be free. (Recall Barrett's sentiments, quoted earlier in this seminar.) Indeed, they seem to have lost entirely the very things that mark freedom: spontaneity; the ability to learn from past experience (especially from mistakes); the ability, and the occasional implementation of that ability, to change or modify behavior; etc.

What this latest thesis ultimately comes down to is this: One lacks free will if one never does anything differently. Sometime or other, having free will has to show up in doing something novel. Let me put it this way: spontaneity, novelty, and the like are not symptoms of having free will; they comprise some of the criteria.¹⁵

¹⁴ Were Kant's daily walks – year in and year out – free choices? I am disinclined to think so. This is not, however, to suggest that they were compulsions. They were probably just a habit, which is not to say that he could not have chosen otherwise. Possibly, they were so much of a routine that he didn't choose at all to walk; he just did it.

¹⁵ I can anticipate a certain objection. In a modern classic, "Time Without Change," Shoemaker (1969) showed how one can have evidence of a passage of time in which there is no change; a fortiori, it is possible that there should be a time with no change. I can imagine someone arguing, along similar lines, that there could be free will in the absence of anyone's ever doing something novel.

Like this would-be objector, I, too, want to invoke Shoemaker's paper and adapt its ingenious technique. For I think it bears out my point, not my opponent's. The analogy must be carried through fully, however.

What Shoemaker's argument showed was not that all eternity could pass without change, but only that some period of time – embedded in a longer time – could pass without change. He showed, that is, that some periods of time could be unaccompanied by change; he did not and – I venture an opinion – could not show that all of time could be unaccompanied by change. And it is just this contrast between 'some period of time' and 'all of history' that I want to analogize with the case of free will. I have no quarrel with it being hypothesized that persons 'have free will' even when they might not be doing something novel; but I do balk at attributing free will to persons none of whom *ever* does anything novel. I am no more prepared to allow the latter than I am to allow that there might be an eternity in which nothing ever changes. In what sense can "eternity" apply to changelessness? In no intelligible sense, I aver.

What has to be true of a world so that someone therein has free will? It seems to me, for the reasons just given, that it will not do to say that someone who always acts in the most regular fashion ‘could’ do something differently. This bare, unrealized *de re* possibility does not distinguish that world from an otherwise identical one in which no one has free will.

For it to be *possible* in world *W* that someone have free will, it is necessary that someone or other in *W* do something novel. (Note the parallelism with the case of there being a river of Coca-Cola. According to the Regularity account, in world *W* there is a physical possibility of there being a river of Coca-Cola only if in *W* there is at some time or other and at some place or other a river of Coca-Cola.) For there to be free will in a world, it is not enough that someone *could do* something different from what he usually does, if he wanted to; it must be that somebody wants to and *does*.

Of course, it is logically impossible that anyone should both do something and do that very thing differently. And it is a vacuous claim to say that somebody does something differently from what he might have done. I mean that someone does something differently from what somebody else, or he himself on a different occasion, has done, is doing, or will do. Really, my point can be put pretty straightforwardly: The existence of freedom requires that there *be* novelty in a world.

If the existence of free will is to depend on the existence of novelty, we can rightly demand an explication of the concept of novelty. What does ‘doing something novel (/new)’ *mean*?

Since I am trying to make having free will a metaphysical matter, it cannot be made to depend on psychological criteria of unexpectedness, for example. Certainly, an event may be ‘novel’ in that it was unexpected; but that someone does something unexpected need not be evidence of her having exercised her free will; the unexpectedness may reflect nothing more than our inattention to what was going on and to what might reasonably have been forecast.

Can ‘novelty,’ understood to be a metaphysical category, even occur in a deterministic world? Can an event be both determined *and* novel?

The concept of novelty appropriate for our purposes has already been illustrated in our earlier case of my choosing which shoes to purchase. Recall that we supposed that I had an unbroken history of always buying brown shoes. Were this behavior to continue, the universal generalization, U_1 , under which my shoe purchases would fall would be: “Whenever ... (uniquely individuating description of me) ... in situation *GHI* buys shoes, he buys brown shoes.” But suppose that tomorrow I do something ‘novel’: I buy black shoes. Thus U_1 cannot be the physical law covering tomorrow’s purchase or, for that matter, any of my previous purchases. The description *GHI*, because of my ‘novel’ behavior tomorrow, turns out to be too broad,

too inclusive. We must add some *differentia*, X , to the specification. *Two* laws – instead of the false U_1 – it would now appear, cover my shoe purchasing behavior: one, L_1 , whose antecedent condition contains the term “ GHI and X ” and whose consequent is “buys brown shoes”; and the other, L_2 , whose antecedent contains the term “ GHI and $\sim X$ ” and whose consequent is “buys black shoes.” By my breaking a pattern of behavior, I have made the world *less uniform* (in the sense previously explained) than it would have been.

Novelty, clearly, can occur in a fully deterministic world. It consists in lessening the uniformity of that world from what it would have been.

In part, in considerable part, what makes the actual world such an *interesting* place to live is that we ourselves, and our fellows, are forever doing new, novel things (i.e., are forever doing things that lessen the uniformity of the world).

A belief in the existence of free will is not threatened by a belief in determinism, that is, the belief that events may be subsumed under physical laws. What does pose a threat to our belief in free will is the idea that determinism carries in its train the doctrine of limited variety,¹⁶ or excessive uniformity, or, putting this still another way, the doctrine that there are ultimately only a few physical laws. What we need, to have free will, is not that our actions should not be subsumable under physical law, but rather that there should be a very great number of physical laws. The freer we are, the more physical laws there are; the less free we are, the fewer physical laws there are.

A world in which there is too much uniformity is a world uncongenial to freedom; so too is a world with too little. But both worlds, at opposite ends of the scale of uniformity, may be fully deterministic worlds. Determinism per se poses no threat to freedom. Certain kinds of deterministic worlds are inhospitable to freedom; but not all.

All too often, discussions about the possibility of free will concentrate attention on the impediments to freedom, for example, on being drugged, hypnotized, mentally ill, physically constrained, etc. But a look at the positive necessary conditions rather than the inhibiting

¹⁶ Discussions of uniformity, variety, and novelty seem to be somewhat out of fashion of late. Some philosophers, e.g., Mill, Keynes, and Carnap, used to talk about these concepts, but these discussions no longer seem topical. They should be. Clarification of these concepts is fundamental to the understanding of the concepts of physical law, determinism, forecastability, and the like.

conditions reveals just how fragile freedom is: The existence and exercise of one's freedom requires a very special kind of world.

What are some of these positive necessary conditions? I hardly know them all; I do not even know whether I know most. But I can specify as least some.

First of all, the world must be reasonably uniform. One must be able to learn from past experiences. One must be able reasonably to anticipate the probable consequences of one's actions, and, unless one can do this a priori (and that seems pretty unlikely for any world fairly similar to this), one must depend on past experience as a guide to the future.

One must be able to make valid inferences pretty regularly. One must be able to select relevant data. One must have a fairly good memory. And one must be able to imagine alternative (possible but not undertaken) courses of action.

Persons who have argued that freedom is impossible if determinism *is* true have probably thought that determinism implies a stifling uniformity. Those who have argued that freedom is impossible if determinism is *not* true have doubtless seen that freedom requires that there be appreciable uniformity and have probably thought that the existence of this degree of uniformity requires that the world be deterministic.

Both sides are correct in their respective views about uniformity: Too much uniformity is inimical to freedom; too little is likewise. And both are incorrect as regards the relation between uniformity and determinism.

Determinism is compatible with both high and low uniformity. Whether a world is congenial to freedom does not depend on that world's being determined, but on the degree of uniformity obtaining in that world.

That Determinism be *false* is not a necessary condition for the existence of free will. But neither does the existence of free will require that Determinism be *true*. Free will exists both in possible worlds that are fully determined and in possible worlds that are not. Only in those worlds in which few or no events are physically determined, that is, in which there is little or no uniformity whatever (i.e., a truly chaotic world), and in those worlds in which all or most events are eternally repeated, that is, in which uniformity is close to or at its maximum, is there no possibility for free will. In short, that free will exists is logically independent of the truth of the Principle of Determinism.