

Chapter One: Physics and Chance

1) Chance

Suppose that the world consisted entirely of point masses, moving in perfect accord with the Newtonian law of motion, under the influence of some particular collection of inter-particle forces. And imagine that that particular law, in combination with those particular forces, allowed for the existence of relatively stable, extended, rigid, macroscopic arrangements of those point masses - chairs (say) and tables and rocks and trees and all of the rest of the furniture of our everyday macroscopic experience.¹ And consider a rock, traveling at constant velocity, through an otherwise empty infinite space, in a world like that. And note that nothing whatsoever in the Newtonian law of motion, together with the laws of the interparticle forces, together with a stipulation to the effect that those interparticle forces are all the forces there are, is going to stand in the way of that rock's suddenly ejecting one of its trillions of elementary particulate constituents at enormous speed and careening off in an altogether different direction, or (for that matter) spontaneously disassembling itself into statuettes of the British royal family, or (come to think of it) reciting the Gettysburg address.

It goes without saying that none of these is in fact a serious possibility. And so the business of producing a scientific account of anything at all of what we actually know of the behaviors of rocks, or (for that matter) of planets or pendula or tops or levers or any of the traditional staples of Newtonian mechanics, is going to call for something over and above the deterministic law of motion, and the laws of the inter-particle forces, and a stipulation to the effect that those inter-particle forces are all the forces there are - something along the lines of a probability-distribution over microconditions, something that will entail, in conjunction with the law of motion and the laws of the inter-particle forces and a stipulation to the effect that those forces are all the forces there are, that the preposterous scenarios mentioned above - although they are not impossible - are nonetheless immensely unlikely.

And there is a much more general point here, a point which has nothing much to do with the ontological commitments or dynamical peculiarities or empirical inadequacies of the mechanics of Newtonian point masses, which goes more or less like this: Take any fundamental physical account of the world on which a rock is to be understood as an arrangement, or as an excitation, or as some more general collective upshot of the behaviors of an enormous number of elementary microscopic physical degrees of freedom. And suppose that there is some convex and continuously infinite set of distinct exact possible microconditions of the world - call that set {R} - each of which is compatible with the macrodescription "a rock of such and such a mass and such and such a shape is traveling at such and such a velocity through an otherwise empty infinite space". And suppose that the fundamental law of the evolutions of those exact microconditions in time is completely deterministic. And suppose that the fundamental law of the evolutions of those exact microconditions in time entails that for any two times $t_1 < t_2$, the values of all of the fundamental physical degrees of freedom at t_2 are invariably some continuous function of the values of those degrees of freedom at t_1 . If all that is the case, then it gets hard to imagine how {R} could possibly fail to include a continuous infinity of distinct conditions in which the values of the elementary microscopic degrees of freedom happen to be lined up with one another in precisely such a way as to produce more or less any preposterous behavior you like - so long as the behavior in question is in accord with the basic ontology of the world, and

with the conservation laws, and with the continuity of the final conditions as a function of the initial ones, and so on.. And so the business of discounting such behaviors as implausible - the business (that is) of underwriting the most basic and general and indispensable convictions with which you and I make our way about in the world - is again going to call for something over and above the fundamental deterministic law of motion, something along the lines, again, of a probability-distribution over microconditions.

If the fundamental microscopic dynamical laws themselves have chances in them, then (of course) all bets are off. But there are going to be chances - or that (at any rate) is what the above considerations suggest - at one point or another. Chances are apparently not to be avoided. An empirically adequate account of a world even remotely like ours in which nothing along the lines of a fundamental probability ever makes an appearance is apparently out of the question. And questions of precisely where and precisely how and in precisely what form such probabilities enter into nature are apparently going to need to be reckoned with in any serviceable account of the fundamental structure of the world.

2) The Case of Thermodynamics

Let's see what there is to work with.

The one relatively clear and concrete and systematic example we have of a fundamental probability-distribution over microconditions being put to useful scientific work is the one that comes up in the statistical-mechanical account of the laws of thermodynamics.

One of the monumental achievements of the physics of the 19th century was the discovery of a simple and beautiful and breathtakingly concise summary of the behaviors of the temperatures and pressures and volumes and densities of macroscopic material systems. The name of that summary is thermodynamics - and thermodynamics consists, in its entirety, of two simple laws. The first of those laws is a relatively straightforward translation into thermodynamic language of the conservation of energy. And the second one, the famous one, is a stipulation to the effect that a certain definite function of the temperatures and pressures and volumes and densities of macroscopic material systems - something called the entropy - can never decrease as time goes forwards. And it turns out that this second law in and of itself amounts to a complete account of the inexhaustible infinity of superficially distinct time-asymmetries of what you might call ordinary macroscopic physical processes. It turns out - and this is something genuinely astonishing - that this second law in and of itself entails that smoke spontaneously spreads out from and never spontaneously collects into cigarettes, and that ice spontaneously melts and never spontaneously freezes in warm rooms, and that soup spontaneously cools and never spontaneously heats up in a cool room, and that chairs spontaneously slow down but never spontaneously speed up when they are sliding along floors, and that eggs can hit a rock and break but never jump off the rock and re-assemble themselves, and so on, without end.

In the latter part of the 19th century, physicists like Ludwig Boltzmann in Vienna and John Willard Gibbs in New Haven began to think about the relationship between thermodynamics and the underlying complete microscopic science of elementary constituents of the entirety of the world - which was presumed (at the time) to be Newtonian Mechanics. And the upshot of those investigations is a beautiful new science called statistical mechanics.

Statistical mechanics begins with a postulate to the effect that a certain very

natural-looking measure on the set of possible exact microconditions of any classical-mechanical system is to be treated or regarded or understood or put to work - of this hesitation more later - as a probability-distribution over those microconditions. The measure in question here is (as a matter of fact) the simplest imaginable measure on the set of possible exact microconditions of whatever system it is one happens to be dealing with, the standard Lebegue measure on the phase-space of the possible exact positions and momenta of the Newtonian particles that make that system up. And the thrust of all of the beautiful and ingenious arguments of Boltzmann and Gibbs, and of their various followers and collaborators, was to make it plausible that the following is true:

Consider a true thermodynamical law, any true thermodynamical law, to the effect that macrocondition A evolves - under such-and-such external circumstances and over such-and-such a temporal interval - into macrocondition B. Whenever such a law holds, the overwhelming majority of the volume of the region of phase-space associated with macrocondition A - on the above measure, the simple measure, the standard measure, of volume in phase-space - is taken up by microconditions which are sitting on deterministic Newtonian trajectories which pass, under the allotted circumstances, at the end of the allotted interval, through the region of the phase space associated with the macrocondition B.

And if these arguments succeed, and if Newtonian mechanics is true, then the above-mentioned probability-distribution over microconditions will underwrite great swaths of our empirical experience of the world: It will entail (for example) that a half-melted block of ice alone in the middle of a sealed average terrestrial room is overwhelmingly likely to be still more melted towards the future, and that a half-dispersed puff of smoke alone in a sealed average terrestrial room is overwhelmingly likely to be still more dispersed towards the future, and that a tepid bowl of soup alone in a sealed average terrestrial room is overwhelmingly likely to get still cooler towards the future, and that a slightly yellowed newspaper alone in a sealed average terrestrial room is overwhelmingly likely to get still more yellow towards the future, and uncountably infinite extensions and variations of these, and incomprehensibly more besides.

But there is a famous trouble with all this, which is that all of the above-mentioned arguments work just as well in reverse, that all of the above-mentioned arguments work just as well (that is) at making it plausible that (for example) the half-melted block of ice I just mentioned was more melted towards the past as well. And we are as sure as we are of anything that that's not right.

And the canonical method of patching that trouble up is to supplement the dynamical equations of motion and the statistical postulate with a new and explicitly non-time-reversal-symmetric fundamental law of nature, a (so-called) past-hypothesis, to the effect that the universe had some particular, simple, compact, symmetric, cosmologically sensible, very low entropy initial macrocondition. The patched-up picture, then, consists of the complete deterministic microdynamical laws and a postulate to the effect that the distribution of probabilities over all of the possible exact initial microconditions of the world is uniform, with respect to the Lebaguse measure, over those possible microconditions of the universe which are compatible with the initial macrocondition specified in the past-hypothesis, and zero elsewhere. And with that amended picture in place, the arguments of Boltzmann and Gibbs will make it plausible not only that that paper will be yellower and ice cubes more melted and people more aged

and smoke more dispersed in the future, but that they were all less so (just as our experience tells us) in the past. With that additional stipulation in place (to put it another way) the arguments of Boltzmann and Gibbs will make it plausible that the second law of thermodynamics remains in force all the way from the end of the world back to its beginning.

*

What we have from Boltzmann and Gibbs, then, is a probability-distribution over possible initial microconditions of the world which - when combined with the exact deterministic microscopic equations of motion - apparently makes good empirical predictions about the values of the thermodynamic parameters of macroscopic systems. And there is a question about what to make of that success: We might take that success merely as evidence of the utility of that probability-distribution as an instrument for the particular purpose of predicting the values of those particular parameters, or we might take that success as evidence that the probability-distribution in question is literally true.

And note (and this is something to pause over) that if the probability-distribution in question were literally true, and if the exact deterministic microscopic equations of motion were literally true, then that probability-distribution, combined with those equations of motion, would necessarily amount not merely to an account of the behaviors of the thermodynamic parameters macroscopic systems, but to the complete scientific theory of the universe - because the two of them together assign a unique and determinate probability-value to every formulable proposition about the exact microscopic physical condition of whatever physical things there may happen to be. If the probability-distribution and the equations of motion in question here are regarded not merely as instruments or inference-tickets but as claims about the world, then there turns out not to be any physical question whatever on which they are jointly agnostic. If the probability-distribution and the equations of motion in question here are regarded not merely as instruments or inference-tickets but as claims about the world, then they are either false or they are in some sense (of which more in a minute) all the science there can ever be.

And precisely the same thing will manifestly apply to any probability-distribution over the possible exact microscopic initial conditions of the world, combined with any complete set of laws of the time-evolutions of those macroconditions.²

*

I want to look into the possibility that the probability-distribution we have from Boltzmann and Gibbs, or something like it, something more up-to date, something adjusted to the ontology of quantum field theory or quantum string theory or quantum brane theory, is true.

And this is a large undertaking.

Let's start slow.

Here are three prosaic observations.

The laws of thermodynamics are not quite true. If you look closely enough, you will find that the temperatures and pressures and volumes of macroscopic physical systems occasionally fluctuate away from their thermodynamically predicted values. And it turns out that precisely the same probability-distribution over the possible microconditions of such a system that accounts so well for the overwhelming reliability of the laws of thermodynamics accounts for the relative frequencies of the various different possible transgressions against those

laws as well. And it turns out that the particular features of that distribution that play a pivotal role in accounting for the overwhelming reliability of the laws of thermodynamics are largely distinct from the particular features of that distribution that play a pivotal role in accounting for the relative frequencies of the various possible transgressions against those laws. It turns out (that is) that the relative frequencies of the transgressions give us information about a different aspect of the underlying microscopic probability-distribution (if there is one) than the overwhelming reliability of the laws of thermodynamics does, and it turns out that both of them are separately confirmatory of the empirical rightness of the distribution as a whole.

And consider a speck of ordinary dust, large enough to be visible with the aid of a powerful magnifying glass. If you suspend a speck like that in the atmosphere, and you watch it closely, you can see it jerking very slightly, very erratically, from side to side, under the impact of collisions with individual molecules of air. And if you carefully keep tabs on a large number of such specks, you can put together a comprehensive statistical picture of the sorts of jerks they undergo - as a function (say) of the temperatures and pressures of the gasses in which they are suspended. And it turns out (again) that precisely the same probability-distribution over the possible microconditions of such a system that accounts so well for the overwhelming reliability of the laws of thermodynamics accounts for the statistics of those jerks too. And it turns out (again) that the particular features of that distribution that play a pivotal role in accounting for the reliability of the laws of thermodynamics are largely distinct from the particular features of that distribution that play a pivotal role in accounting for the statistics of the jerks. And so the statistics of the jerks give us information about yet another aspect of the underlying microscopic probability-distribution (if there is one), and that new information turns out to be confirmatory, yet again, of the empirical rightness of the distribution as a whole.

And very much the same is true of isolated pin-balls balanced atop pins, or isolated pencils balanced on their points. The statistics of the directions in which such things eventually fall turn out to be very well described by precisely the same probability-distribution over possible microconditions, and it turns out (once more) that the particular features of that distribution that play a pivotal role in accounting for the reliability of the laws of thermodynamics are distinct from the particular features of that distribution that play a pivotal role in accounting for the statistics those fallings.

And so the standard statistical posit of Boltzmann and Gibbs - when combined with the microscopic equations of motion - apparently has in it not only the thermodynamical science of melting, but also the quasi-thermodynamical science of chance fluctuations away from normal thermodynamic behavior, and (on top of that) the quasi-mechanical science of unbalancing, of breaking the deadlock, of pulling infinitesimally harder this way or that. And these sorts of things are manifestly going to have tens of thousands of other immediate applications. And it can now begin to seem plausible that this standard statistical posit might in fact have in it the entirety of what we mean when we speak of anything's happening at random or just by coincidence or for no particular reason.

3) The Special Sciences in General

The thermodynamic parameters all have straightforward and explicit and complete translations into the fundamental physical languages of Newtonian point mechanics, or of non-relativistic quantum mechanics, or of relativistic quantum field theory, or what have you.

But that's not the case, and perhaps it will never be the case, and perhaps it can never be the case, for (say) economics, or epidemiology, or reader reception theory.

Let's take that for granted, then. Let's suppose that there can be no explicit translations from the languages of the various special sciences into the language of fundamental physics. But let's suppose as well that there is some fundamental physical language in which the world can be described completely, in at least the minimal sense that no two physically possible worlds can have different descriptions in the languages of any of the special sciences unless they have different descriptions in that fundamental physical language too. And let's imagine that we have in hand a complete microscopic dynamics, a complete theory (that is) of the time-evolutions of the elementary particles and fields or the elementary strings and branes or the elementary quantum-mechanical wave-functions or whatever the elementary physical constituents of the world turn out to be. And let's imagine (just for the sake of keeping things simple, and just for the moment) that that theory is fully deterministic.

And now take a computer, or (rather) super-computer, or (rather) a super-duper-computer, or whatever sort of a computer it is that the operations about to be described might turn out to require. And enter the dynamics into the computer. And enter the exact microscopic physical conditions of (say) the entirety of the solar system, at precisely 8:00 P.M., on the evening of a certain particular formal dinner party, into the computer.⁵ And instruct the computer to perform a calculation (which nothing on the level of principle will now stand in the way of its successfully carrying out, with the benefit of no further input whatever, and to whatever accuracy we might like) of the fundamental physical conditions obtaining near the surface of the earth throughout the period between (say) 10:00 PM and 10:15 PM later that same evening. And instruct the computer to output that information in such a way as to make it possible for its human operators to walk at will, in real time, about a virtual reconstruction of the barometric and electromagnetic conditions in the room where the dinner-party is taking place, throughout the interval between 10:00 P.M. and 10:15 P.M. later that same evening, and look here, and listen there, and just sort of take the whole thing in.

A sufficiently powerful computer (then) equipped with nothing over and above the fundamental physical laws, and provided (on the particular occasion in question) with nothing over and above an appropriate set of fundamental physical initial conditions, can show us, can display, for us, how any particular party (or war, or election, or painting, or investigation, or marriage, or whatever) comes out.

Suppose we want something fancier. Suppose (that is) that we want a computer that can do more than merely show us or display for us how the dinner party in question (or any other one) is going to come out. Suppose that what we want is a computer that can evaluate for us, and report to us, and predict for us, in English - equipped (mind you) with nothing over and above the fundamental physical theory, and provided with nothing over and above an appropriate set of fundamental physical initial conditions (of which more in a minute) - whether the party succeeds or fails.

Here's how to do that: Take a computer of the sort we were talking about before. And enter the dynamics into the computer. And enter the exact microscopic physical conditions of the entirety of the solar system, at precisely 8:00 P.M., on the evening of the party, into the computer. And instruct the computer (as before) to perform a calculation of the fundamental physical conditions obtaining near the surface of the earth throughout the period between 10:00 PM and 10:15 PM later that same evening. And instruct the computer to output that information in precisely the same sort of virtual-reality format as we described above. And instruct the

computer's human operator to survey whatever portions of that output he needs to in order to evaluate whether the party turns out to be a success or a failure, and to report his findings in writing, in English, on a piece of paper, and to place that piece of paper in a certain particular (otherwise empty) box. And now enter the dynamics into a second computer. And enter into that second computer the exact microscopic fundamental physical conditions of the entirety of (say) the sealed interstellar spaceship containing the first computer and its output mechanisms and its human operator and the piece of paper and the empty box and whatever else happens to be in there at (say) the moment just after all of the instructions described in the earlier part of this paragraph have been delivered. And instruct the second computer to calculate and reproduce for us the physical contents of the box twenty minutes or so hence.

The output of this second computer, then, will consist of a piece of paper on which either "the party is a success" or "the party is a failure" appears. And the report on that piece of paper will (with extremely rare exceptions, of which more in a minute) prove accurate - once the actual historical facts are in. And the input (once again) consists of nothing over and above the fundamental physical laws and the fundamental physical initial conditions of various parts of the world back at the time the party in question first got underway.

Suppose we want something fancier still. Suppose we want a computer that can deduce for us, given nothing over and above the fundamental physical laws, the laws (if there are any) of the special science of the success or failure of formal dinner parties. Suppose (for example) that we want a computer that can deduce for us, given nothing over and above the fundamental physical laws, whether inviting an odd number of guests, or inviting an even number, is more likely to produce a better party.

Here's how to work that: Take a computer. And enter the fundamental physical dynamical laws into it. And set things up in such a way as to allow the operator of the computer to take a virtual tour - of just the sort we talked about above - of whichever of the possible exact physical microconditions of the solar system he chooses.⁴ And instruct the operator to survey the space of those exact physical microconditions by means of this technique, and to identify the regions of that space which correspond to circumstances in which odd numbers of guests are being invited to a formal dinner party, and to identify the regions of that space which correspond to circumstances in which even numbers of guests are being invited to a formal dinner party, and to order the computer to calculate, and to display for him, how all of those dinner-parties come out, and to prepare a written report, and to put it into the box we were talking about above. And now (just as before) enter the dynamics into a second computer. And enter into that second computer the exact microscopic fundamental physical conditions of the entirety of the sealed interstellar spaceship containing the first computer and its output mechanisms and its human operator and the piece of paper and the empty box and whatever else happens to be in there at the moment just after all of the instructions described in the earlier part of this paragraph have been delivered. And instruct the second computer to calculate and reproduce for us the physical contents of the box twenty minutes (or twenty years, or however long it might imaginably take) hence.

Now a few remarks are in order.

Note that the roles of the various computers in the above three scenarios are merely to function as concrete realizations of the implicative structures of the fundamental physical theories - the roles of those computers are merely to make vivid what sorts of information about the world those theories have in them, what sorts of things they can predict, what sorts of things

they can account for. Nothing whatever hinges on the possibility (which is presumably immensely remote - both for practical reasons and for reasons of principle as well) of such computers ever actually being constructed!

And note that the role of the actual living human operator of the computer in the second and third of the scenarios is merely to serve as a catalogue of fundamental physical initial conditions. Neither of those two scenarios involves any living human being's ever actually evaluating the success or failure of a virtual dinner party, or surveying sets of possible initial conditions, or preparing a written report, or anything of the sort - all of that gets done by a simulator, by a subroutine - call it the 'human operator subroutine' - which exists only in cyberspace, and of which the living human being in question is merely the plan, merely the template, merely the set of instructions. The living human being in question can perfectly well be far away, or asleep, or dead, long before the simulated evaluations or surveyings or reportings ever get underway. It is nothing whatever over and above the computer, nothing (that is) over and above the implicative structure of the fundamental physical laws, that's doing all the work.

These subroutines (by the way) - precisely because and precisely in so far as they can faithfully reproduce the behavioral dispositions of the living human beings on which they are modeled, beings which (after all) can lie and make mistakes and get moody and have heart attacks - will necessarily be imperfect as speakers of the language of dinner parties. But they can be very good. They can be exactly as good as the best of us, or as whole committees of us, or as whole societies of us. And the very idea of doing any better, the very idea of coming up with a program which somehow instantiates a formal and mechanical and algorithmic scheme for translating from the language of fundamental physical theory to the language of dinner parties is (by hypothesis) out of the question.

And there's one more thing - and this is precisely the thing that the previous two sections of this chapter were about. Recall the third of the scenarios we talked about above - the one where the computer derives the general laws (if - once again - there are any) of the success and failure of dinner parties, and the full apparatus for explaining why this or that particular dinner party succeeded or failed, from nothing over and above the fundamental laws of physics. The computer first surveys the space of possible exact physical microconditions of the solar system, and then it identifies the regions of that space which correspond to circumstances in which odd numbers of guests are being invited to a formal dinner party and the regions of that space which correspond to circumstances in which even numbers of guests are being invited to a formal dinner party, and then it calculates how all those dinner parties come out, and finally it prepares a written report on the respective probabilities that even or odd numbers of invitations will (as a matter of fundamental physical law) result in success. And what I want to draw attention to at the moment is just that the fundamental physical laws in question here are going to have to include stipulations of a kind that we have so far neglected to bring up in this section, that the fundamental physical laws in question here are going to have to amount (again) to something over and above the complete theory of the time-evolutions of the fundamental physical systems that we talked about some pages back. And the reason (of course) is this: A complete theory of the time-evolutions of the fundamental physical systems - although it will settle all questions as to which initial microconditions lead to success and which to failure, although it will settle all questions as to how many of the microconditions which correspond to (say) an even number of invitations being sent out lead to success and how many lead to failure - will have nothing whatever to say about the probability of any particular one of those microconditions actually obtaining, given (say) that an even number of invitations are sent out. And without that we have

nothing of any macroscopic use at all. And so (again) we are going to need an other and altogether different sort of fundamental physical law, a law which will apparently need to take the form of a probability-distribution - or something like a probability-distribution - over initial conditions.

And the alluring possibility is (again) that the law we need here is in fact the one we already have, the one suggested by altogether different sorts of considerations, the one that seemed to show some promise of making concrete and explicit and quantitative sense of what we mean when we speak of anything's happening at random or just by coincidence or for no particular reason, the one (that is) due to Boltzmann and Gibbs.

4) Explanation

Here is a line of argument - one of many - aimed directly against the sort of universality and completeness of physics that I was trying to imagine in the previous section. It comes from Science, Truth, and Democracy, by my friend and teacher Philip Kitcher. The worry here is not about the capacities of fundamental physical theories to predict - which Philip is willing to grant - but about their capacities to explain. Philip directs our attention to

“.... the regularity discovered by John Arbuthnot in the early eighteenth century. Scrutinizing the record of births in London during the previous 82 years, Arbuthnot found that in each year a preponderance of the children born had been boys: in his terms, each year was a “male year”. Why does this regularity hold? Proponents of the Unity-of Science view can offer a recipe for the explanation, although they can't give the details. Start with the first year (1623); elaborate the physicochemical details of the first copulation-followed-by-pregnancy showing how it resulted in a child of a particular sex; continue in the same fashion for each pertinent pregnancy; add up the totals for male births and female births and compute the difference. It has now been shown why the first year was “male”; continue for all subsequent years. Even if we had this “explanation” to hand, and could assimilate all the details, it would still not advance our understanding. For it would not show that Arbuthnot's regularity was anything more than a gigantic coincidence. By contrast, we can already give a satisfying explanation by appealing to an insight of R. A. Fisher. Fischer recognized that, in a population in which sex ratios depart from 1:1 at sexual maturity, there will be a selective advantage to a tendency to produce the underrepresented sex. It will be easy to show from this that there should be a stable evolutionary equilibrium at which the sex ratio at sexual maturity is 1:1. In any species in which one sex is more vulnerable to early mortality than the other, this equilibrium will correspond to a state in which the sex ratio at birth is skewed in favor of the more vulnerable sex. Applying this analysis to our own species, in which boys are more likely than girls to die before reaching puberty, we find that the birth sex ratio ought to be 1.104:1 in favor of males - which is what Arbuthnot and his successors have observed. We now understand *why* [my italics], for a large population, all years are overwhelmingly likely to be male.”

The key word here, the word that carries the whole burden of Philip's argument, is ‘coincidence’.

And that will be worth pausing over, and thinking about.

The moral of the first section of this chapter (remember) was that the fundamental physical laws of the world, merely in order to get the narrowest imaginable construal of their 'work' done, merely in order to get things right (that is) about projectiles and levers and pulleys and tops, will need to include a probability-distribution over possible microscopic initial conditions. And once a distribution like that is in place, all questions of what is and isn't likely, all questions of what was and wasn't to be expected, all questions of whether or not this or that particular collection of events happened merely 'at random' or 'for no particular reason' or 'as a matter of coincidence', are (in principle) settled. And (indeed) it is only by reference to a distribution like that that talk of coincidence can make any precise sort of sense in the first place - it is only against the background of a distribution like that that questions of what is or is not coincidental can even be brought up.

It goes without saying that we do not (typically, consciously, explicitly) consult that sort of a distribution when we are engaged in the practical business of making judgments about what is and is not coincidental. But that is no evidence at all against the hypothesis that such a distribution exists, and it is no evidence at all against the hypothesis that such a distribution is the sole ultimate arbiter of what is and is not coincidental, and it is no evidence at all against the hypothesis that such a distribution informs every single one of our billions of everyday deliberations. If anything along the lines of the complete fundamental theory we have been trying to imagine here is true (after all) some crude, foggy, reflexive, largely unconscious but perfectly serviceable acquaintance with that distribution will have been hard-wired into us as far back as when we were fish, as far back (indeed) as when we were slime, by natural selection - and lies buried at the very heart of the deep instinctive primordial unarticulated feel of the world. If anything along the lines of the complete fundamental theory we have been trying to imagine here is true (after all) the penalty for expecting anything else, the penalty for expecting anything to the contrary, is extinction.

And if one keeps all this in the foreground of one's attention, it gets hard to see what Philip can possibly have in mind in supposing that something can amount to a 'gigantic coincidence' from the standpoint of the true and complete and universal fundamental physical theory of the world and yet (somehow or other) not be.

*

If anything along the lines of the picture we are trying to imagine here should turn out to be true, then any correct special-scientific explanation whatsoever can in principle be uncovered, can in principle be descried, in the fundamental physical theory of the world, by the following procedure:

Start out with a distribution of probabilities which is uniform, on the standard statistical-mechanical measure, over all of the possible exact initial microconditions of the world which are compatible with the past-hypothesis, and zero elsewhere. And conditionalize that distribution on whatever particular features of the world play a role in the special-scientific explanation in question - conditionalize that distribution (that is) on whatever particular features of the world appear either explicitly or implicitly among the explanans of the special-scientific explanation in question⁵. And check to see whether or not the resultant distribution - the conditionalized distribution, makes the explanandum likely. If it does, then we have recovered the special-scientific explanation from the fundamental physical theory - and if it doesn't, then

either the fundamental theory, or the special-scientific explanation, or both, are wrong.

Consider (for example) the evolution of the total entropy of the universe over the past ten minutes. That entropy (we are confident) is unlikely to have gone down over those ten minutes. The intuition is that the entropy's having gone down over those ten minutes would have amounted to a gigantic coincidence. The intuition is that the entropy's having gone down over those ten minutes would have required detailed and precise and inexplicable correlations among the positions and velocities of all of the particles that make the universe up. And questions of whether or not correlations like that are to be expected, questions of whether or not correlations like that amount to a coincidence, are matters (remember) on which the sort of fundamental physical theory we are thinking about here can by no means be agnostic. And it is part and parcel of what it is for that sort of a theory to succeed that it answers those questions correctly. It is part and parcel of what it is for that sort of a theory to succeed (that is) that it transparently captures, and makes simple, and makes elegant, and makes precise, the testimony our intuition, and our empirical experience of the world, to the effect that correlations like that are in fact fantastically unlikely, that they are not at all to be expected, that they do indeed amount to a gigantic coincidence. And there is every reason in the world to believe that there is a fundamental physical theory that can do that. It was precisely the achievement of Boltzmann and Gibbs (after all) to make it plausible that the Newtonian laws of motion, together with the statistical-postulate, together with the past-hypothesis, all of it conditionalized on a proposition to the effect that the world was not swarming, ten minutes ago, with malevolent Maxwellian Demons, can do, precisely, that.

And now consider the descent of man. The first humans (we are confident) are unlikely to have condensed out of swamp gas, or to have grown on trees, or to have been born to an animal incapable of fear. The first (after all) would require detailed and precise and inexplicable correlations among the positions and velocities of all of the molecules of swamp gas, and of the surrounding air, and the ground, and god knows what else. And the second would require a vast, simultaneous, delicately co-ordinated unimaginably fortuitous set of mutations on a single genome. And the third would require that every last one of a great horde of mortal dangers all somehow conspire to avoid the animal in question - with no help whatever from the animal herself - until she is of age to deliver her human child. And it is precisely because the account of the descent of man by random mutation and natural selection involves vastly fewer and more minor and less improbable such coincidences than any of the imaginable others that it strikes us as the best and most plausible explanation of that descent we have. And (indeed) it is precisely the relative paucity of such coincidences, and it is precisely the relative smallness of whatever such coincidences there are, to which words like 'random' and 'natural' are meant to direct our attention! And questions about what does and what does not amount to a coincidence are matters (once again) on which the sort of fundamental physical theory we are imagining here can by no means be agnostic. And it is part and parcel of what it is for that sort of a theory to succeed (once again) that it answers every last one of those questions correctly.

Now, compelling arguments to the effect that this or that particular fundamental physical theory of the world is actually going to be able to do all that are plainly going to be harder to come by here than they were in the much more straightforward case of the entropy of the universe. All we have to go on are small intimations - the ones mentioned above, the ones you can make out in the behaviors of pin-balls and pencils and specks of dust - that perhaps the exact microscopic laws of motion together with the statistical postulate together with the past-hypothesis has in it the entirety of what we mean when we speak of anything's happening at

random or for no particular reason or just by coincidence.

But if all that should somehow happen to pan out, if there is a true and complete and fundamental physical theory of the sort that we have been trying to imagine here, then it is indeed going to follow directly from the fundamental laws of motion, together with the statistical-postulate, together with the past-hypothesis, all of it conditionalized on the existence of our galaxy, and of our solar system, and of the earth, and of life, and of whatever else is implicitly being taken for granted in scientific discussions of the descent of man, that the first humans are indeed extraordinarily unlikely to have condensed out of swamp gas, or to have grown on trees, or to have been born to an animal incapable of fear.

And very much the same sort of thing is going to be true of the regularity discovered by Arubthot.

What Fisher has given us (after all) is an argument to the effect that it would amount to a gigantic coincidence, that it would represent an enormously improbable insensitivity to pressures of natural selection, that it would be something very much akin to a gas spontaneously contracting into one particular corner of its container, for sex ratios to do anything other than settle into precisely the stable evolutionary equilibrium that he identifies. And questions about what does and what does not amount to a coincidence are (for the last time) matters on which the sort of fundamental physical theory we are imagining here can by no means be agnostic. And it is part and parcel of what it is for that sort of a theory to succeed that it answers every last one of those questions correctly.

And once again, compelling arguments to the effect that this or that particular fundamental physical theory of the world is actually going to be able to do all that are plainly going to be hard to come by - and all we are going to have to go on are the small promising intimations from pin-balls and pencils and specks of dust.

But consider how things would stand if all that should somehow happen to pan out. Consider how things would stand if there is a true and complete and fundamental physical theory of the sort that we have been trying to imagine here.

Start out - as the fundamental theory instructs us to do - with a distribution of probabilities which is uniform, on the standard statistical-mechanical measure, over all of the possible exact initial microconditions of the world which are compatible with the past-hypothesis, and zero elsewhere. And evolve that distribution - using the exact microscopic deterministic equations of motion - up to the stroke of midnight on December 31st of 1623. And conditionalize that evolved distribution on the existence of our galaxy, and of our solar system, and of the earth, and life, and of the human species, and of cities, and of whatever else is implicitly being taken for granted in any scientific discussion of the relative birth rates of boys and girls in London in the years following 1623. And call that evolved and conditionalized distribution P_{1623} .

If there is a true and complete and fundamental theory of the sort that we have been trying to imagine here, then what Fisher has given will amount to an argument that P_{1623} is indeed going to count it as likely that the preponderance of the babies born in London, to human parents, in each of the 82 years following 1623, will be boys. Period. End of story.

Of course, the business of explicitly calculating P_{1623} from the microscopic laws of motion and the statistical postulate and the past-hypothesis is plainly, permanently, out of the question. But Philip's point was that even if that calculation could be performed, even (as he says) "if we had this "explanation" to hand, and could assimilate all the details, it would still not advance our understanding. For it would not show that Arbuthnot's regularity was anything

more than a gigantic coincidence.” And this seems just....wrong. And what it misses - I think - is that the fundamental physical laws of the world, merely in order to get the narrowest imaginable construal of their ‘work’ done, merely in order to get things right (that is) about projectiles and levers and pulleys and tops, are going to have to come equipped, from the word go, with chances.

*

And those chances are going to bring with them - in principle - the complete explanatory apparatus of the special sciences. And more than that: those chances, together with the exact microscopic equations of motion, are going to explain all sorts of things about which all of the special sciences taken together can have nothing whatever to say, they are going to provide us - in principle - with an account of where those sciences come from, and of how they hang together, of how it is that certain particular sets of too-ings and fro-ings of the fundamental constituents of the world can simultaneously instantiate every last one of them, of how each of them separately applies to the world in such a way as to accommodate the fact that the world is a unity.

And so (you see) what gets in the way of explaining things is not at all the conception of science as unified, but the conceit that it can somehow not be.

5) The General Business of Legislating Initial Conditions

All of this delicately hangs (of course) on the possibility of making clear metaphysical sense of the assignment of real physical chances to initial conditions⁶.

And conceptions of chance as anything along the lines of (I don’t know) a cause or a pressure or a tendency or a propensity or a pulling or a nudging or an enticing or a cajoling or (more generally) as anything essentially bound up with the way in which instantaneous states of the world succeed one another in time, is manifestly not going to be up to the job - since the initial condition of the world is (after all) not the temporal successor of anything, and there was (by definition) no historical episode of the world’s having been pulled or pressed or nudged or cajoled into this or that particular way of getting started.

Our business here (then) is going to require another understanding of chance. And an understanding of law in general, I think, to go with it. Something Humean. Something wrapped up not with an image of governance, but with an idea of description. Something (as a matter of fact) of the sort that’s been worked out, with slow and sure and graceful deliberation, over these past twenty years or so, by David Lewis and Barry Loewer.

*

Here’s the idea. You get to have an audience with God. And God promises to tell you whatever you’d like to know. And you ask Him to tell you about the world. And He begins to recite the facts: such-and-such a property (the presence of a particle, say, or some particular value of some particular field) is instantiated at such-and-such a spatial location at such-and-such a time, and such-and-such another property is instantiated at such-and-such another spatial

location at such-and-such another time, and so on. And it begins to look as if all this is likely to drag on for a while. And you explain to God that you're actually a bit pressed for time, that this is not all you have to do today, that you are not going to be in a position to hear out the whole story. And you ask if maybe there's something meaty and pithy and helpful and informative and short that He might be able to tell you about the world which (you understand) would not amount to everything, or nearly everything, but would nonetheless still somehow amount to a lot. Something that will serve you well, or reasonably well, or as well as possible, in making your way about in the world.

And what it is to be a law, and all it is to be a law, on this picture of Hume's and Lewis' and Loewer's, is to be an element of the best possible response to precisely this request - to be a member (that is) of that set of true propositions about the world which, alone among all of the sets of true propositions about the world that can be put together, best combines simplicity and informativeness.

On a picture like this, the world, considered as a whole, is merely, purely, there. It isn't the sort of thing that is susceptible of being explained or accounted for or traced back to something else. There isn't anything that it obeys. There is nothing to talk about over and above the totality of the concrete particular facts. And science is the business of producing the most compact and informative possible summary of that totality. And the components of that summary are called laws of nature.⁷

The world (on this picture) is not what it is in virtue of the laws being what they are, the laws are what they are in virtue of the world's being what it is.

*

Now, different possible worlds - different possible totalities (that is) of concrete particular facts - may turn out to accommodate qualitatively different sorts of maximally compact and informative summaries.

The world might be such that God says: "I have just the thing: The furniture of the universe consists entirely of particles. And the force exerted by any particle on any other particle is equal to the product of the masses of those two particles divided by the square of the distance between them, directed along the line connecting them. And those are all the forces there are. And everywhere, and at every time, the acceleration of every particle in the world is equal to the total force on that particle at that time divided by its mass. That won't tell you everything. It won't tell you nearly everything. But it will tell you a lot. It will serve you well. And it's the best I can do, it's the most informative I can be, if (as you insist) I keep it short." Worlds like that are called (among other things) Newtonian and particulate and deterministic and non-local and energy-conserving and invariant under Galilean transformations.

Or the world might be different. The world might be such that God says "Look, there turns out not to be anything I can offer you in the way of simple, general, exact, informative, exceptionlessly true propositions. The world turns out not to accommodate propositions like that. Let's try something else. Global physical situations of type A are followed by global physical situations of type B roughly (but not exactly) 70% of the time, and situations of type A are followed by situations of type C roughly (but not exactly) 30% of the time, and there turns out not to be anything else that's simple to say about which particular instances of A-situations are followed by B-situations and which particular instances of the A-situations are followed by C-situations. That's pithy too. Go fourth. It will serve you well." We speak of worlds like

that as being lawful but indeterministic - we speak of them as having real dynamical chances in them.

Or the world might be such that God says: “Sadly, I have nothing whatever of universal scope to offer you - nothing deterministic and nothing chancy either. I’m sorry. But I do have some simple, useful, approximately true rules of thumb about rainbows, and some others about the immune system, and some others about tensile strength, and some others about birds, and some others about interpersonal relationships, and some others about stellar evolution, and so on. It’s not elegant. It’s not all that concise. But it’s all there is. Take it. You’ll be glad, in the long run, that you did.” We speak of worlds like that - following Nancy Cartwright - as dappled.

Or the world might be such that God has nothing useful to offer us at all. We speak of worlds like that as chaotic - we speak of them as radically unfriendly to the scientific enterprise.

Or the world might (finally) be such that God says: “All of the maximally simple and informative propositions that were true of the Newtonian particulate deterministic non-local energy-conserving Galilean-invariant universe are true of this one too. The furniture of the universe consists entirely of particles. And the force exerted by any particle on any other particle is equal to the product of the masses of those two particles divided by the square of the distance between them, directed along the line connecting them. And those are all the forces there are. And everywhere, and at every time, the acceleration of every particle in the world is equal to the total force on that particle at that time divided by its mass. But that’s not all. I have something more to tell you as well. Something (as per your request) simple and helpful and informative. Something about the initial condition of the world. I can’t tell you exactly what that condition was. It’s too complicated. It would take too long. It would violate your stipulations. The best I can do by way of a simple and informative description of that condition is to tell you that it was one of those which is typical with respect to a certain particular probability-distribution - the Boltzmann-Gibbs distribution, for example. The best I can do by way of a simple and informative description of that initial condition is to tell you that it was precisely the sort of condition that you would expect, that it was precisely the sort of condition that you would have been rational to bet on, if the initial condition of the world had in fact been selected by means of a genuinely dynamically chancy procedure where the probability of this or that particular condition’s being selected is precisely the one given in the probability-distribution of Boltzmann and Gibbs.” And this is precisely the world we encounter in classical statistical mechanics. And this is the sought-after technique - or one of them - for making clear metaphysical sense of the assignment of real physical chances to initial conditions. The world has only one microscopic initial condition. Probability-distributions over initial conditions - when they are applicable - are compact and efficient and informative instruments for telling us something about what particular condition that is.⁸

And note that it is of the very essence of this Humean conception of the law that there is nothing whatever metaphysical at stake in the distinctions between deterministic worlds, and chancy ones, and dappled ones, and chaotic ones, and ones of the sort that we encounter in a deterministic statistical mechanics. All of them are nothing whatever over and above totalities of concrete particular facts. They differ only in the particular sorts of compact summaries that they happen - or happen not - to accommodate.

6) Dynamical Chances

Quantum Mechanics has fundamental chances in it.

And it seems at least worth inquiring whether or not those chances can do us any good. It seems worth inquiring (for example) whether or not those chances are up to the business of guaranteeing that we can safely neglect the possibility of a rock, traveling at constant velocity, through an otherwise empty infinite space, spontaneously disassembling itself into statuettes of the British royal family. And the answer turns out to depend, interestingly, sensitively, on which particular one of the available ways of making sense of Quantum Mechanics as a universal theory, on which particular one of the available ways (that is) of solving the quantum-mechanical measurement problem, turns out to be right.

The sorts of chances that come up in orthodox pictures of the foundations of Quantum Mechanics - the pictures (that is) that have come down to us from the likes of Bohr and von Neumann and Wigner - turn out not to be up to the job. On pictures like those, the chanciness that is so famously characteristic of the behaviors of quantum-mechanical systems enters into the world exclusively in connection with the act of measurement. Everything whatever else - according to these pictures - is fully and perfectly deterministic. And there are almost certainly exact microscopic quantum-mechanical wave-functions of the world which are compatible with there being a rock, traveling at constant velocity, through an otherwise empty infinite space, and which are sitting on deterministic Quantum-Mechanical trajectories along which, a bit later on, if no 'acts of measurement' take place in the interim, that rock spontaneously disassembles itself into statuettes of the British royal family. And it happens to be the case, it happens to be an empirical fact, that the overwhelming tendency of rocks like that not to spontaneously disassemble themselves into statuettes of the British royal family has nothing whatsoever to do with whether or not, at the time in question, they are in the process of being measured!

And the same thing goes (for slightly different reasons) for the chances that come up in more precisely formulable and recognizably scientific theories of the collapse of the wave-function like the one due to Penrose. On Penrose's theory, quantum-mechanical chanciness enters into the evolution of the world not on occasions of 'measurement', but (rather) on occasions when certain particular wave-functions of the world - wave-functions corresponding to superpositions of macroscopically different states of the gravitational field - obtain. But the worry here is that there may be exact microscopic quantum-mechanical wave-functions of the world which are compatible with there being a rock, traveling at constant velocity, through an otherwise empty infinite space, and which are sitting on deterministic Quantum-Mechanical trajectories which scrupulously avoid all of the special collapse-inducing macroscopic superpositions mentioned above, and along which, a bit later on, that rock spontaneously disassembles itself into statuettes of the British royal family.

And the same thing goes (for slightly more different reasons) for the chances that come up in Bohm's theory. The only things that turn out to be chancy, on Bohm's theory, are the initial positions of the particles. The only sort of fundamental chance there is in Bohm's theory is (more particularly) the chance that the initial spatial configuration of all of the particles in the world was such-and-such given that the initial quantum-mechanical wave-function of those particles was so-and-so. And it happens - on Bohm's theory - that those parts of the fundamental physical laws that govern the evolution of the wave-function in time, and those parts of the fundamental physical laws that stipulate precisely how the evolving wave-function drags the particles around, are completely deterministic. And it turns out that there are possible exact wave-functions of the world which are compatible with there initially being a rock,

traveling at constant velocity, through an otherwise empty infinite space, which (if those laws are right) will determine, all by themselves, that the probability of that rock's spontaneously disassembling itself into statuettes of the British royal family is overwhelmingly, impossibly, high.

And the long and the short of it is that the same thing goes (for all sorts of different reasons) for the chances that come up in Modal theories, and in the many-worlds interpretation, and in the Ithaca interpretation, and in the transactional picture, and in the relational picture, and in a host of other pictures too.

On every one of those theories, the business of guaranteeing that we can safely neglect the possibility of a rock, traveling at constant velocity, through an otherwise empty infinite space, doing something silly, turns out to require the introduction of another species of chance into the fundamental laws of nature - something over and above and altogether unrelated to the Quantum-Mechanical chances, something (more particularly) along the lines of the non-dynamical un-quantum-mechanical probability-distributions over initial microscopic conditions of the world that we have been discussing throughout the earlier sections of this chapter.

And this seems (I don't know) odd, cluttered, wasteful, sloppy, redundant, perverse.

And there is (perhaps) a way to do better. There is a simple and beautiful and promising theory of the collapse of the quantum-mechanical wave-function due to Ghirardi and Rimini and Weber that puts the quantum-mechanical chanciness in differently.

On the GRW theory - as opposed to (say) Bohm's theory, quantum-mechanical chanciness is dynamical. And on the GRW theory - as opposed to any theory whatever without a collapse of the wave-function in it - quantum-mechanical chanciness turns out to be a chanciness in the time-evolution of the universal wave-function itself. And on the GRW theory - as opposed to theories of the collapse like the one due to Penrose - the intrusion of quantum-mechanical chanciness into the evolution of the wave-function has no trigger; the probability of a collapse per unit time (that is) is fixed, once and for all, by a fundamental constant of nature; the probability of a collapse over the course any particular time-interval (to put it one more way) has nothing whatsoever to do with the physical situation of the world over the course of that interval.

And this is precisely what we want. On the GRW theory - as opposed to any of the other theories mentioned above, or any of the other proposed solutions to the measurement problem of which I am aware - quantum-mechanical chanciness is the sort of thing that there can be no outwitting, and no avoiding, and no shutting off. It insinuates itself everywhere. It intrudes on everything. It seems fit (at last) for all of the jobs we have heretofore needed to assign to probability distributions over initial conditions. If the fundamental dynamics of the world has this sort of chanciness in it, then there will be no microconditions whatsoever - not merely very few, not merely a set of measure zero, but not so much as a single one - which make it likely that a rock, traveling at constant velocity, through an otherwise empty infinite space, will spontaneously disassemble itself into statuettes of the British royal family⁹. And the same thing presumably goes for violations of the second law of thermodynamics, and for violations of the law of the survival of the fittest, and for violations of the law of supply and demand.

And so if something along the lines of the GRW theory should actually turn out to be true, science will apparently be in a position to get along without any probability-distribution whatsoever over possible initial microconditions.¹⁰ If something along the lines of the GRW theory should actually turn out to be true, then it might imaginably turn out that there is at

bottom only a single species of chance in nature. It might imaginably turn out (that is) that all of the robust lawlike statistical regularities there are in the world are at bottom nothing more or less than the probabilities of certain particular GRW collapses hitting certain particular sub-atomic particles.¹¹

Whether or not it does turn out to be true - of course - is a matter for empirical investigation.

1. And this, of course, is not true. And it is precisely because Newtonian Mechanics appears not to allow for the existence of these sorts of things or, even for the stability of the very atoms that make them up, that it is no longer entertained as a candidate for the fundamental theory of the world. But put all that aside for the moment.

2. Shelly Goldstein and Detlef Durr and Nino Zanghi and Tim Maudlin have worried, with formidable eloquence and incisiveness, that probability-distributions over the initial conditions of the world might amount to vastly more information than we could ever imaginably have a legitimate epistemic right to. Once we have a dynamics (once again) a probability-distribution over the possible exact initial conditions of the world will assign a perfectly definite probability to the proposition that I am sitting precisely here writing precisely this precisely now, and to the proposition that I am doing so not now but (instead) 78.2 seconds from now, and to the proposition that the Yankees will win the world series in 2097, and to the proposition that the zodiac killer was Mary Tyler Moore, and to every well-formed proposition whatever about the physical history of the world. And it will do so as a matter of fundamental physical law. And the worry is that it may be mad to think that there could be a fundamental physical law as specific as that, or that we could ever have good reason to believe anything as specific as that, or that we could ever have good reason to believe anything that logically implies anything as specific as that, even if the calculations involved in spelling such an implication out are prohibitively difficult.

Moreover, there are almost certainly an enormous number of very different probability-distributions over the possible initial conditions of the world which are capable of underwriting the laws of thermodynamics more or less as well as the standard, uniform, Boltzmann-Gibbs distribution does. And the reasons for that will be worth rehearsing in some detail.

Call the initial macrocondition of the world M . And let R_M be that region of the exact microscopic phase-space of the world which corresponds to M . And let aR_M be the sub-region of R_M which is taken up with “abnormal” microconditions - microconditions (that is) that lead to anomalously widespread violations of the laws of thermodynamics. Now, what the arguments of Boltzmann and Gibbs suggest is as a matter of fact not only that the familiarly calculated

volume of aR_M is overwhelmingly small compared with the familiarly calculated volume of R_M - which is what I have been at pains to emphasize so far - but also that aR_M is scattered, in unimaginably tiny clusters, more or less at random, all over R_M . And so the percentage of the familiarly calculated volume of any regularly shaped and not unimaginably tiny sub-region of which is taken up with abnormal microconditions will be (to an extremely good approximation) the same as the percentage of the familiarly calculated volume of R_M as a whole which is taken up by aR_M . And so any reasonably smooth probability distribution over the microconditions in R_M - any probability-distribution over the microconditions in R_M (that is) that varies slowly over distances two or three orders of magnitude larger than the diameters of the unimaginably tiny clusters of which aR_M is composed - will yield (to an extremely good approximation) the same overall statistical propensity to thermodynamic behavior as does the standard uniform Boltzmann-Gibbs distribution over R_M as a whole. And exactly the same thing, or much the same thing, or something in the neighborhood of the same thing, is plausibly true of the behaviors of pin-balls and adrenal glands and economic systems and everything else as well.

The suggestion (then) is that we proceed as follows: Consider the complete set of those probability-distributions over the possible exact initial conditions of the world - call it $\{P_f\}$ - which can be obtained from the uniform Boltzmann-Gibbs distribution over R_M by multiplication by any relatively smooth and well-behaved and appropriately normalized function f of position in phase space. And formulate your fundamental physical theory of the world in such a way as to commit it to the truth of all those propositions on which every single one of the probability-distributions in $\{P_f\}$, combined with the dynamical laws, agree - and to leave it resolutely agnostic on everything else.

If everything works as planned, and if everything in the paragraph before last is true - a theory like that will entail that the probability of smoke spreading out in a room, at the usual rate, is very high, and it will entail that the probability of a fair and well flipped coin's landing on heads is very nearly 1/2, and it will entail (more generally) that all of the stipulations of the special sciences are very nearly true. And yet (and this is what's different, and this is what's cool) it will almost entirely abstain from the assignment of probabilities to universal initial conditions. It will entail - and it had better entail - that the probability that the initial condition of the universe was one of those that lead to anomalously widespread violations of the laws of thermodynamics, that the probability that the initial condition of the universe lies (that is) in aR_M , is overwhelmingly small. But it is going to assign no probabilities whatsoever to any of the smoothly-bounded or regularly-shaped or easily-describable proper subsets of the microconditions compatible with M .

Whether or not a theory like that is ever going to look as simple and as serviceable and as perspicuous as the picture we have from Boltzmann and Gibbs (on the other hand) is harder to say. And (anyway) I suspect that at the end of the day it is not going to spare us the awkwardness of assigning of a definite probability, as a matter of fundamental physical law, to the proposition that the Zodiac Killer was Mary Tyler Moore. I suspect (that is) that every single one of the probability-distributions over R_M that suffice to underwrite the special sciences are going end up assigning very much the same definite probability to the proposition that the Zodiac Killer was Mary Tyler Moore as the standard, uniform, Boltzmann-Gibbs distribution does. And if that's true, then a move like the one being contemplated here may end up buying us very little.

And beyond that, I'm not sure what to say. In so far as I can tell, our present business is

going proceed in very much the same way, and arrive at very much the same conclusions, whether it starts out with the standard, uniform, Boltzmann-Gibbs probability-distribution over the microconditions in R_M , or with any other particular one of the probability-distributions in $\{P_f\}$, or with $\{P_f\}$ as a whole. And the first of those seems by far the easiest and the most familiar and the most intuitive and the most explanatory and (I guess) the most advisable. Or it does at first glance. It does for the time being. It does unless, or until, we find it gets us into trouble.

3. Or (at any rate) enter in the conditions of as much of the world, at 8:00 P.M. of the evening in question, as would need to be taken into account in the process of calculating the physical conditions on the surface of the earth some hours later - enter in (that is) the conditions throughout the cross-section, at 8:00 P.M. of the evening in question, of the backward light cone of the surface of the earth some hours later.

4. Things might be set up in such a way as to allow the operator to point and click on a map of the space of states.

5. Those explanans, of course, are initially going to be given to us in the language of one or another of the special sciences. And so, in order to carry out the sort of conditionalization we have in mind here, we are going to need to know which of those special-scientific explanans correspond to which regions of the space of possible exact physical microconditions of the world. And those correspondences can be worked out - not perfectly, mind you, but to any degree of accuracy and reliability we like - by means of the super-duper computational techniques described in section 3.

6. I will be taking it for granted here that a probability-distribution over initial conditions, whatever else it is, is an empirical hypothesis about the way the world contingently happens to be.

But this is by no means the received view of the matter. Indeed, the statistical postulate of Boltzmann and Gibbs seems to have been understood by its inventors as encapsulating something along the lines of an a priori principle of reason, a principle (more particularly) of indifference, which runs something like this: Suppose that the entirety of what you happen to know of a certain system S is that S is X . And let $\{v_i\}_{X,t}$ be the set those of the possible exact microconditions of S such that v_i 's obtaining at t is compatible with S 's being X . Then the principle stipulates that for any two $v_j, v_k \in \{v_i\}_{X,t}$ the probability of v_j 's obtaining at t is equal to the probability of v_k 's obtaining at t .

And that (I think) is more or less what the statistical postulate still amounts to in the imaginations of many physicists. And that (to be sure) has a supremely innocent ring to it. It sounds very much, when you first hear it, as if it is instructing you to do nothing more than attend very carefully to what you mean, to what you are saying, when you say that the entirety of what you know of S is that S is X . It sounds very much as if it is doing nothing more than reminding you that what you are saying when you say something like that is that S is X , and (moreover) that for any two $v_j, v_k \in \{v_i\}_{X,t}$, you have no more reason for believing that v_j obtains at t than you have for believing that v_k obtains at t , that (in so far as you know) nothing favors any particular one of the $v_j \in \{v_i\}_{X,t}$ over any particular other one of the $v_j \in \{v_i\}_{X,t}$,

that (in other words) the probability of any particular one of those microconditions obtaining at t , given the information you have, is equal to the probability of any particular other one of them obtaining at t .

And this is importantly and spectacularly wrong. And the reasons why it's wrong (of which there are two: a technical one and a more fundamental and less often remarked-upon one too) are worth rehearsing.

The technical reason has to do with the fact that the sort of information we can actually have about physical systems - the sort that we can get (that is) by measuring - is invariably compatible with a continuous infinity of the system's microstates. And so the only way of assigning equal probability to all of those states at the time in question will be by assigning each and every one of them the probability zero. And that will of course tell us nothing whatever about how to make our predictions.

And so people took to doing something else - something that looked to them to be very much in the same spirit - instead. They abandoned the idea of assigning probabilities to individual microstates, and took instead to stipulating that the probability assigned to any finite region of the phase space which is entirely compatible with X - under the epistemic circumstances described above - ought to be proportional to the continuous measure of the points within that region.

But there's a trouble with that - or at any rate there's a trouble with the thought that it's innocent - too. The trouble is that there are in general an infinity of equally mathematically legitimate ways of putting measures on infinite sets of points. Think, for example, of the points on the real number line between 0 and 1. There is a way of putting measures on that set of points according to which the measure of the set of points between any two numbers a and b (with $a < 1$ and $b < 1$ and $b > a$) is $b - a$, and there is another way of putting measures on that set of points according to which the measure of the set of points between any two numbers a and b between (with $a < 1$ and $b < 1$ and $b > a$) is $b^2 - a^2$, and according to the first of those two formulae there are "as many" points between 1 and $\frac{1}{2}$ as there are between $\frac{1}{2}$ and 0, and according to the second of those two formulae there are three times "as many" points between 1 and $\frac{1}{2}$ as there are between $\frac{1}{2}$ and 0, and there turns out to be no way whatever (or at any rate none that anybody has yet dreamed up) of arguing that either one of these two formulae represents a truer or more reasonable or more compelling measure of the "number" or the "amount" or the "quantity" of points between a and b than the other one does. And there are (moreover) an infinite number of other such possible measures on this interval as well, and this sort of thing (as I mentioned above) is a very general phenomenon.

And anyway, there is a more fundamental problem, which is that the sorts of probabilities being imagined here, probabilities (that is) conjured out of airy nothing, out of pure ignorance, whatever else might be good or bad about them, are obviously and scandalously unfit for the sort of explanatory work that we require of the probabilities of Boltzmann and Gibbs. Forget (then) about all the stuff in the last three paragraphs. Suppose there was no trouble about the measures. Suppose that there were some unique and natural and well-defined way of expressing, by means of a distribution-function, the fact that "nothing in our epistemic situation favors any particular one of the microstates compatible with X over any other particular one of them". So what? Can anybody seriously think that that would somehow explain the fact that the actual microscopic conditions of actual thermodynamic systems are statistically distributed in the way that they are? Can anybody seriously think that it is somehow necessary, that it is

somehow a priori, that the particles of which the material world is made up must arrange themselves in accord with what we know, with what we happen to have looked into? Can anybody seriously think that our merely being ignorant of the exact microstates of thermodynamic systems plays some part in bringing it about, in making it the case, that (say) milk dissolves in coffee? How could that be? What can all those guys have been up to? If probabilities have anything whatever to do with how things actually fall out in the world (after all) then knowing nothing whatever about a certain system other than X can in and of itself entail nothing whatever about the relative probabilities of that system's being in one or another of the microstates compatible with X; and if probabilities have nothing whatever to do with how things actually fall out in the world, then they can patently play no role whatever in explaining the behaviors of actual physical systems; and that would seem to be all the options there are to choose from!

7. This is not at all (of course) to deny that there are such things as scientific explanations! There are all sorts of explanatory relations - on a picture like this one - among the concrete particular facts, and (more frequently) among sets of the concrete particular facts. There are all sorts of things to be said (for example) about how smaller and more local patterns among those facts fit into, or are subsumed under, or are logically necessitated by, larger and more universal ones. But the totality of the concrete particular facts is the point at which - on a view like this one - all explaining necessarily comes to an end.

8. The strategy described in footnote 2 - the strategy (that is) of abstaining from the assignment of any particular probability-distribution over those of the possible microconditions of the world which are compatible with its initial macrocondition, has sometimes been presented as a way around the problem, as a way of avoiding the problem, of making clear metaphysical sense of assigning probability-distributions to the initial conditions of the world. But that seems all wrong - for two completely independent reasons. First, the strategy in question makes what looks to me to be ineliminable use of sets of probability-distributions over the possible initial microconditions of the world - and if those distributions themselves can't be made sense of, then (I take it) sets of them can't be made sense of either. Second, the problem of making clear metaphysical sense of the assignment of probability-distributions to the initial microcondition of the world isn't the sort of thing that needs getting around - since (as we have just now been discussing) it can be solved!

9. For details, arguments, clarifications, and any other cognitive requirements to which this sentence may have given birth - see chapter seven of my Time and Chance.

10. It will still be necessary (mind you) to include among the fundamental laws of the world a stipulation to the effect that the world started out in some particular low-entropy macrocondition - but (in the event that something along the lines of GRW should turn out to be true) nothing further, nothing chancy, nothing (that is) along the lines of a probability-distribution over those of the possible microconditions of the world which are compatible with that macrocondition, will be required.

These considerations are spelled out in a great deal more detail in chapter seven of my Time and Chance.

11. The theory we are envisioning here will of course assign no probabilities whatever to possible initial microconditions of the world, and it will consequently assign no perfectly definite probabilities to any of the world's possible conditions - microscopic or otherwise - at any time in its history. What it's going to do - instead - is to assign a perfectly definite probability to every proposition about the physical history of the world given that the initial microcondition of the world was A, and another perfectly definite probability to every proposition about the physical history of the world given that the initial microcondition of the world was B, and so on. But note that the probability that a theory like this is going to assign to any proposition P given that the initial microcondition of the world was A is plausibly going to be very very very very close to the probability that it assigns to P given that the initial microcondition of the world was B - so long as both A and B are compatible with the world's initial macrocondition, and so long as P refers to a time more than (I don't know) a few milliseconds into the world's history.