

TRANSITION CHANCES AND CAUSATION

Frank Arntzenius
Dept. of Philosophy
University of Southern California
Los Angeles, CA 90089-0451

ABSTRACT

The general claims of this paper are as follows. As a result of chaotic dynamics we can usually not know what the deterministic causes of events are. There will however be invariant forwards transition chances from earlier types of events, which we call the causes, to later types of events, which we call the effects. There will be no invariant backwards transition chances between these types of events. This asymmetry has the same origin and explanation as the arrow of time of thermodynamics.

TRANSITION CHANCES AND CAUSATION

1. Introduction

Consider a generic type epidemic of some kind of disease, let us call it disease X. One observes that initially, prior to the epidemic, the transition frequency, from being healthy to acquiring the symptoms of disease X is low, and that at the height of the epidemic, it becomes relatively high. Epidemeologists would immediately start to search for a cause, or causes of this disease.

One of the ways in which they would do this is that they would start to search for a factor, or factors, Y, such that the transition frequencies (per unit of time, e.g. per week) from having property Y and being healthy to having disease X are roughly constant in time. They could for instance find that although the frequency of transition from being healthy to acquiring X varies strongly in time (as the epidemic follows its course) the transition frequency from being-healthy-and-being-in-close-contact-with-someone-who-has-X to having-X (in the next week) is roughly constant in time and relatively high, while the transition frequency from not-being-in-contact to having-X (in the next week) is roughly constant in time and very low. They would then call it a contagious disease, and say that at least one of the causes of the disease is contact with a diseased person. Of course, they would then soon advise people to stay away from people that have disease X. If this advice were followed strictly the epidemic would soon die down (though not disappear completely, since the transition frequency from not-having-contact to having-X, though low, is not 0).

As the epidemic dies down they also notice that there is a peculiar time-asymmetry in the transition frequencies: the forwards transition frequencies are invariant but the backwards ones are not. That this is so is most easily seen if we assume that the advice of the epidemeologists is eventually followed very strictly. For when that advice is adhered to by almost everyone, there will be a very low incidence of the disease, and almost all the cases of the disease will then have been cases in which the disease was not preceded by contact with a diseased person. On the other hand

at the height of the epidemic almost all cases would have been preceded by contact with a diseased person. Thus the backwards transition frequency from having-X to having-contact-with-someone-with-X (in the previous week) will be high before the advice is followed, but low after the advice is followed.

More generally, one can show that for typical processes for which one has time-invariant non-zero transition frequencies from a set of properties A_i to another set of properties B_j , and the frequencies of properties A_i vary (significantly) in time, i.e. one is not in (approximate) equilibrium, then the transition frequencies from B_j to A_i must vary (significantly) in time. See the Appendix, part II, for details. Thus, in the generic case, one will have non-invariant backwards transition frequencies if one has invariant forwards transition frequencies. Let me now give another example to illustrate this asymmetry and its connection with causation.

Let us assume that in the year 1750 AD there were very few smokers in Europe, say 10 out of every 100 adults, as is reasonable to suppose since they only recently discovered tobacco. Suppose that about 50% of those smokers acquired lungcancer, and that only 10 % of the non-smokers acquired lungcancer. It then follows that the majority, indeed about 2/3, of lungcancer sufferers were not previously smokers.

Suppose that in 1950 about 50 out of every 100 adults smoked. Let us suppose that percentage of smokers that later acquires lungcancer is still the same: namely about 50%, and let us assume that the percentage of non-smokers that later acquired lungcancer is also still the same, namely 10%. It then follows that the majority of lungcancer sufferers, 5 out of every 6, previously was a smoker.

What we see here is that the forwards transition frequencies are invariant in time, and the backwards transition frequencies are highly non-invariant. What this strongly suggests is that there are invariant forwards transition chances from being a smoker to acquiring lung cancer, and that there are either no such backwards transition chances, or that they are non-invariant in time. And thus it suggests that smoking determines your chances of later getting lungcancer, but that having

lungcancer does not determine your chances of being a smoker at an earlier time, and thus that smoking causes lungcancer, but lungcancer does not cause smoking. More generally I claim that in nature there are invariant forwards transition chances and frequencies, and no such invariant backwards transition chances or frequencies. Moreover, when we ask why an event of a particular type occurred, what we are often searching for is a set of factors such that the transition chances from these factors to events of that type are invariant in time. We typically also make further demands, namely that the transition chances be invariant under further conditionalisation, that they be Markovian, and that the common cause principle be satisfied. However, for the most part I will ignore those demands in this paper. Now let me turn to another example in order to explain the connection with statistical mechanics and chaos.

2. Statistical mechanics and the transition asymmetry

Suppose a friend gives you a thousand identical toy robots which, if you wind them up and put them down, every second take a step to the left or right, or stay where they are. Suppose you have a circle divided into 20 segments, numbered 1 through 20 (segment 20 borders segment 1), where each segment has exactly the size of the steps that the robots can take. Suppose that you initially, at time $t=0$, put all of the robots down in the same segment, say segment 5. You then let them wander around the circle. You notice that each toy makes each step about equally often, so that the toys after a while become more or less evenly distributed on the circle and remain more or less evenly distributed for as long as you observe them. As you watch this you realize that there is a peculiar time-asymmetry to the developments you have been watching. When you look at any time at any particular spot you find that one second later almost always approximately one third of the toys that were at that spot have moved one segment to the right, one third have moved one segment to the left, and one third have remained in place. However, the same is not true in the other direction of time. If you look, say, at spot 6 at time $t=1$ you will find that all of them were at spot 5 at the earlier time $t=0$. But if you look at that same spot 6 at a much later time, say time 184, you

will find that approximately $1/3$ of those were at spot 5 at the earlier time 183 (the others coming from segments 6 and 7 at $t=183$). That is to say: the relative proportion of toys at spot 6 which, 1 second earlier, were at spot 5, varies in time. Indeed in general you notice that the backwards transition frequencies change in time while the forwards ones do not.

If one now would want to say that there are transition chances governing this process, then the most plausible conclusion is that the steps of the toys, when left alone, are governed by position independent and time independent forwards transition chances (namely chance $1/3$ for each step each second), and not by position and time independent backwards transition chances, and thus that the laws governing this process are time-asymmetric. Another way of putting this is as follows: if one runs a film of this process in the forwards direction it appears that at all times it is governed by the same forwards transition chances. But if one runs the film in the backwards direction it appears initially as if the transition chances everywhere are $1/3$ for each step, but then towards the end probabilities of steps towards point 0 become bigger and bigger.

Now let me briefly discuss how one could explain this asymmetry in the phenomena, assuming that the laws of development of microphysical states which underlie these phenomena are time-reversible and deterministic. Suppose that in fact the robot decides which way to go in the following manner. It has a hollow body with a gas consisting of many molecules in it. One of these molecules is radioactive. The robot has three geiger counters in it which monitor in which of 3 areas the radioactive atom is, a millisecond before it makes the step. If it is in the top one then it will make a step to the right, if it is in the middle one it makes no step, if it is in the bottom one it makes a step to the left. The molecule bounces around like all molecules of a gas, and the time between steps is sufficiently large and its motion is sufficiently chaotic such that its next position is not correlated to its previous position and hence consecutive steps are not correlated.

Now let us ask: what explains the fact that the forwards transition probabilities are always $1/3$. The obvious answer is: because, due to the molecular chaos, at each time, and for each position the robot could be in, the chance of the radioactive molecule being in any one of the 3 compartments

is $1/3$. Since this determines each time which step is taken this explains why the chance of each step always is $1/3$. Thus it seems that we have a very good explanation of the invariance of the forwards transition probabilities, on the basis of the assumption of molecular chaos.

However, there is a problem with this explanation. For, the same argument can be run to show that the backwards transition probabilities should be invariant, and we know that that can not be right. Let me give you the analogous argument.

Assuming that the robot is an isolated deterministic system, just as the microscopic state of the robot just prior to a step determines what step it will take, the microscopic state of the robot just after it has made a step will also determine what the step was that it previously took: determinism works in both directions of time, at least the determinism of classical physics does. In fact, the position of the radioactive molecule $1/1000.000$ of a second after the step has been taken almost completely determines what step was taken $1/1000.000$ second earlier. For in that very short time it most likely has not moved from one compartment to another, so that almost always it correctly indicates what the step was that just was taken. Hence this position almost always reliably indicates what the previous position of the robot was. Now, if the positions of the radioactive molecule, just after a step has been taken, would also have chance $1/3$ of being in each compartment, no matter what time it is and what position the robot occupies, then one could similarly infer that each backwards transition always has chance $1/3$, and that the backwards transition frequencies should always approximately equal $1/3$. But we know that that can't be right: they are not always $1/3$.

In fact, let us look at the microscopic states of the robots just before and after the first step. Before the first step all the robots are in position $x=0$, and the radioactive molecules are equally likely to be in each compartment. Just after the first step, the robots are distributed equally over $x=-1$, $x=0$ and $x=1$. But the radioactive molecules are not equally likely to be in each compartment. For the robots with the molecule in the bottom compartment all moved to the left, the robots with the molecules in the middle one all stayed where they were, and the ones with the molecule in the top all moved to the right, and these molecules haven't had the time yet to move. Thus at $t=1.00000001$

we have an unequal distribution of microscopic states for each macroscopic position of the robots. This explains why the backwards steps are not equally likely at each position. By time 1.9999999, due to molecular chaos, the radioactive molecules will again be more or less equally distributed over the compartments, and hence the forward step at that time will be equally likely. But at time 2.0000001 we will have another unequal distribution across the compartments for the different robot positions, which will explain the backwards transition frequencies that obtain at time $t=2$.

What do we conclude from this? We conclude that we can explain the invariance of the forwards transition probabilities by assuming that the microscopic conditions which together with the macroscopic state determine the forwards transitions are always equally likely, or uniformly distributed, whereas the microscopic conditions which together with the macroscopic state determine the backwards transitions are not always equally likely, i.e. uniformly distributed.

Now this may strike you as a crazy assumption to make and not much of an explanation of anything. And I think you indeed have grounds to complain and be mystified. All I want to point out at this point is that the assumption that we made, namely that the probability of the forwards development of each macroscopic state, and not the probability of the backwards development of a macroscopic state, should be computed by assuming a uniform probability distribution of microscopic states compatible with that macroscopic state, is a basic assumption made in statistical mechanics in order to explain many thermodynamic time irreversible phenomena. That is to say: the asymmetry in the phenomena that I have noted, the approximate invariance of forwards transition frequencies, is, if one assumes underlying determinism, a consequence of standard assumptions made in statistical mechanics.

In order to make the relation of the transition chance asymmetry to the asymmetries of statistical mechanics even clearer let me indicate how it relates to the best known asymmetry of statistical mechanics: the development towards thermodynamic equilibrium in the forwards direction of time. Let us again consider the toy example. Take any initial chance distribution of positions of the toys on the circle. No matter what it is like, it will always develop monotonically towards the

uniform, equilibrium, distribution, in which all positions on the circle are equally likely. And any initial frequency distribution will most likely develop towards the uniform equilibrium distribution. More generally, in part I of the Appendix I outline some well known results in the theory of Markov processes which show that any initial distribution whose development is governed by invariant forwards transition chances, such that either there are finitely many possible states, or the transition chances satisfy a very mild constraint called "condition D", must develop towards an equilibrium distribution. Hence, for such processes, any initial frequency distribution is very likely to develop into an approximate equilibrium distribution.

In part II of the appendix I show that whenever the forwards transition chances, and the initial distribution are such that this distribution will develop towards an equilibrium state, then the backwards transition chances can not be invariant, unless the initial distribution was already an equilibrium distribution. Thus in every case that one has an asymmetric development towards equilibrium one can not have both invariant forwards and backwards transition chances. Thus there is a strong link indeed between the asymmetry of invariant forwards transition chances and non-invariant backwards transition chances and the thermodynamic asymmetry of development towards an equilibrium forwards in time.

Next let me briefly discuss the ways in which people have attempted to justify the assumption that the microscopic states compatible with a given macroscopic state will be uniformly distributed. The standard way in which people have attempted to do so is by arguing that the development of microscopic states is so chaotic, that even if they were not uniformly distributed at some time, they would be pretty soon afterwards, and would remain so during any further chaotic development.

This idea has enormous intuitive appeal. Suppose you start with a deck of cards in which all the red cards are separated from the black cards, and you then shuffle them in a random (chaotic) manner for a few minutes. Surely, it is very likely that after a while the reds and blacks will be distributed more or less evenly through the pack, and surely the chances of finding a black at any

particular spot will soon become 50%. Thus as long as the development of a set of states is chaotic enough, it would appear that the distribution of the states will soon become uniform.

A further appealing feature of this approach is that it promises to explain why the difference between macroscopic and microscopic properties is made. The general idea is that one can look at the laws of development, the dynamics, of the systems in question, in order to see what properties develop in a chaotic manner, and which ones do not. The chaotically developing ones can be expected at any time to be distributed in a uniform manner, while the non-chaotically developing will not necessarily be so distributed. There will be patterns of development, transition chances, between the non-chaotically developing states, which will be time-invariant, due to the time-invariance of the distribution of the chaotically developing states. Indeed it may not always be the case that it is the microscopic (small scale) properties that are developing chaotically, and the macroscopic (large scale) properties that are not. The laws of dynamics will entail which properties will be uniformly distributed and which not.

Given the intuitive appeal, and interesting consequences of this approach, much effort has been made by physicists in the 20th century, in order to examine the extent to which the development of actual systems and idealized systems are chaotic, in various senses of the term "chaotic". And, of course, physicists and philosophers have spent much time and effort discussing the merits and implications of the results produced by this research. This is not the place to examine all these results and debates, or for a definitive judgments of the merits of each of these results. But let me briefly mention some of these, and suggest ways in which they might relate to the transition asymmetry that I have indicated and causal relations.

One has arrived at a classification of various ways in which the development of systems can be chaotic: ergodic, mixing, Bernoulli etc. And of various systems plus laws of development, mostly idealized ones, it has been shown that they are chaotic in these various senses. These notions of chaotic dynamics are all entirely time-symmetric: if they are chaotic in the forward direction of time so are they in the backwards direction of time. However, as I have said before, given that the

dynamics is chaotic one standardly assumes, for purposes of forming expectations about future states, that the chaotically developing factors will distribute themselves uniformly, while one does not, and could not consistently, assume the same regarding backwards developments.

One of the nice features of the formal results in dynamics is that the dynamics alone (assuming a given division into measure 0 and non-measure 0 sets in statespace) uniquely picks out a particular uniform distribution of microscopic state, a unique measure on the statespace, as the way in which the microscopic states will distribute themselves. And experimentally, when it has been tested, the forwards transition chances have exactly as one would expect on such a picture. There appears thus to be a strong relation between chaos and uniformity. For the predictions made in statistical mechanics on the assumption that one has uniform probability distributions w.r.t. forwards developments when one has a sufficient amount of chaos, have been very exact and very well confirmed. How much the accuracy of these predictions depends on the chaotic nature of the developments is a matter of debate, but there is no doubt that extremely accurate predictions have been made with respect to forwards developments, on the assumption of uniform (microcanonical, canonical, etc.) distributions when there is a sufficient amount of chaos, and that similar "predictions" with respect to backwards developments in non-equilibrium situations are completely wrong. Thus it would seem that in nature chaos is related to invariant forwards transition chances and not to invariant backwards transition chances, and that we would do well to expect invariant forwards transition chances when we encounter chaos.

The more vague and general picture that emerges from the considerations so far is as follows. Factors that develop in a totally chaotic manner, loose all memory of their past states. Such factors can be assumed to have chances of being in states that are specified by the uniform equilibrium distribution, when it comes to calculating the chances of future developments. They are, as it were, "distributed uniformly w.r.t future developments". Factors that do not develop in a chaotic manner do retain a (complete, or more incomplete) "memory" of their past states. For deterministic systems the development of such factors will depend on the current state of that factor and some other non-

chaotically developing factors, and the current state of some other factors that have been developing in a chaotic manner. Therefore such factors will develop according to invariant forwards transition chances. Although theoretical results for idealized systems can give us some idea as to which factors develop in a non-chaotic manner, and which develop in a chaotic manner, in practice we uncover this by searching for such invariant forwards transition chances. Now let us consider 2 more examples which will allow me to make my account slightly more precise, and slightly amend it.

3. Pegboards and strings in the wind

First let us consider a case that one can easily reproduce at home. Get a ball and flat piece of wood that is much bigger than the ball. Attach a large number of pegs to the wood, all more or less at equal distances from each other, where the distances are slightly larger than the diameter of the ball. Then attach wooden edges to the wood, and angle the whole construction at such an angle that if you let the ball go at the top, it will gently role down while bouncing off many pegs. Then get 200 of your friends to do exactly the same. Now pick 2 small regions at the top, A and B, and divide the bottom into 20 segments. Then at your command, 150 of you should let the balls go from A, and 50 of you should let it go from B, all at the same time. In each case note in which segment the ball arrives. Next, you should repeat this procedure, except that now 150 should let it go from B, and 50 from A. You will find that the forwards transition frequencies from A are approximately the same in run 1 and in run 2, and that the forwards transition frequencies from B are also approximately the same in run 1 and in run 2. I.e. the forwards transition frequencies are invariant in time. Indeed, if the board is long enough (compared to its width), and there are enough pegs, then the ones that left from A will distribute themselves in approximately the same way as the ones that left from B, in both runs. But the backwards transition frequencies will not be invariant in time. For instance consider the case where we do not have a long enough board to get identical transition frequencies from A and B. Then consider a region, 11, such that about 20% of those departing from A arrive there, and about 10% of those departing from B arrive there. Then in the first run about 30 balls

departing from A will arrive at 11, and about 5 balls from B. So the backwards transition frequency from 11 to B is about 1 in 7, ie. about 14% in run 1. In run 2 though, about 10 balls departing from A will arrive at region 11, and about 15 from B. Thus in run 2 the backwards transition frequency from 11 to B is about 60%, quite different. What does all of this have to do with chaos?

Well, consider all the trajectories originating from region A. Some balls released from region A will pass the first peg they encounter to the left and others will pass it to the right. Indeed, each peg that can be reached from A will partition and "reshuffle" the set of trajectories originating from region A into 2 further subdivisions. Consequently a uniform initial distribution of positions in A will arrive as a distribution that was chopped up into very small bits, and shuffled many times, at the bottom of the pegboard¹. This will result in some distribution over the coarse partition of the final state space that corresponds to the division of the final positions into 20 segments. And one would expect that the actual frequency distributions over the 20 final segments, which one observes when one looks at the final segments that balls arrive in when one lets them go from A, should, almost always, be close to this distribution. Why would one expect this? After all, a completely uniform distribution is a very special distribution. Different people, all told to release all the balls from within region A, can surely not all be adequately represented as having exactly equal probability (density) for releasing them from any point within region A.

Well, consider some other, but smooth, initial distribution over A. This distribution will be chopped up and shuffled in exactly the same way as the uniform one was. This will mean that the final tiny "chops" will now be of different heights, and each chop might not be entirely flat ("on top"). However, because of the thorough shuffling, the taller "chops" will occur almost equally often in each of the final segments, and the same will be true of chops of any approximate height. But if that is so, then, even though the final distribution is quite a different one, the probabilities of each of the 20 final segments will be approximately the same as it was when we assumed uniform initial distribution. I.e. the final distribution in the coarse grained state space corresponding to the final 20 segments will be roughly the same. Indeed, even if the initial distribution in A was "choppy", rather

than smooth, as long as that "choppiness" is not correlated to the way in which the parts of the initial distribution get redistributed over the final segments, the final distribution over the coarse state space, corresponding to the division into 20 final segments, will be almost exactly the same. For instance, suppose that the initial choppy distribution was such that its tiny chops occupy exactly and only the tiny choppy initial areas within A that will develop into region 11. Then, indeed the final distribution will attribute probability 1 to region 11, and probability 0 to all other regions. But only if there is such a correlation between the initial distribution and the initial chops that correspond to arrivals in particular regions, will such an initial distribution not result in approximately the same distribution over the final segments. Of course, if one releases all balls from region A, and there is no such initial correlation, it follows that the time reverse of the final distribution will incorporate such a correlation between final chops and the region that each chop will develop into.

So, in order to explain the invariance of the forwards transition chances, one needs only assume the lack of correlations between initial microscopic positions and the detailed structure of the chaotic reshuffling, rather than a complete uniformity of the initial distribution of microscopic position compatible with the macroscopic positions.

Before turning to my next example, let me examine whether there is not an inconsistency lurking in my argument, as informal as it is. I have argued that, due to the chaos in the development of positions, most initial fine-grained distributions in a given region A will develop into roughly the same final coarse distribution over segments 1-20. And the same is true of most microscopic distributions over initial regions B, and for most initial distributions over C etc. . That is to say, microscopic distributions over different regions B, C, etc. might lead to different final coarse distributions, but most microscopic distributions confined to the same initial region will lead to approximately the same coarse final distribution. And that in turn implies that most initial fine-grained distributions, not confined to any particular initial region, will develop such that the forwards transition chances from each initial coarse grained area A, B,.... to each final coarse grained area 1,2,... will be almost exactly the same.

But similarly, most microscopic distributions confined to, say, region 11, would develop backwards in time (or if you wish, would develop forwards in time, if you time reversed each state in such a distribution) into roughly the same coarse-grained distribution over regions A, B,..... . And that implies that most microscopic distributions over the arrival regions would be such that the coarse-grained backwards transition chances from regions 1, 2,.... to regions A, B,.... will be almost exactly the same. But this appears to be inconsistent with my previous claim that most initial microscopic distributions imply the near invariance of the coarse grained forwards transition chances. For, to each initial microscopic initial distribution there corresponds a unique microscopic final distributions. So what is true of most microscopic initial distributions must also be true of most microscopic final distributions. But we know that (near) invariant forwards transition chances are typically not compatible with (near) invariant backwards transition chances. (See Part II of the appendix for details.) So how could both be the case for most microscopic distributions?

The resolution to this worry is that the (near) invariance of the coarse grained forwards transition chances and the (near) invariance of the coarse grained backwards transition chances are compatible when we have an (approximate) equilibrium coarse-grained distribution. In this case the (approximate) coarse-grained equilibrium distribution is the (approximately) uniform distribution that attributes (approximately) equal chances to the equal sized regions A, B,... and (approximately) equal chances to the equal sized regions 1, 2,

Now, I claim that the reason there is no inconsistency is that the vast majority of microscopic initial, or final, distributions will correspond to an approximately uniform coarse-grained initial, and final, distribution. Of course, this claim relies on some measure that we put on the space of all possible distributions. But, one can make such a claim plausible by considering how many ways there are of distributing a huge amount of small parts (to be thought of as the parts of the possible microscopic distributions) into 20 sets. Assuming equal likelihood of each part in each set, one will overwhelmingly likely end up with an approximately even distribution. Thus most microscopic distributions correspond to both invariant forwards and backwards transition chances in coarse

grained space. Now of course, this argument relied on some intuitive measure or probability assignment. But without any such intuitive measure one could say nothing at all.

Ultimately, the confirmation and justification of such plausibility arguments is experiment. This particular plausibility argument shows, at the very least, that my previous argument as to when to expect invariant forwards transition chances is coherent. It does not explain why there is such an asymmetry of initial and final conditions, just that it can be simply and coherently stated, and related to chaotic dynamics. Moreover, as far as I can tell, experience and experiment tells us that it seems to be empirically correct too.

For our final example let us consider a test of the resilience of a particular type of string to enormous gusts of wind. Suppose we want to know how likely it is that certain strings will break in high winds. To find this out we take 1000 strings, which, as far as we can tell, are identical, put them in 1000 wind tunnels, and expose each one of them to, as far as we can tell, identical enormous gusts of wind at regular intervals, say 1000 times, each gust 1 minute apart from the next. We note, at the end of each individual gust which strings have broken and which ones have not. We keep applying the gusts, even if the strings have already previously broken.

Now, whether a string on any given occasion will break will presumably depend not only on its observable macroscopic structure (the type of string its thickness etc.) and the observable structure of the gusts (their approximate strength, direction etc.) but presumably also on its microscopic structure and the microscopic structure of the gusts. We might expect that the prior development that leads to the microscopic nature of the windgusts and the changes in the microscopic structure of the strings before during and after windgusts is pretty chaotic, and hence we might expect invariant forwards transition chances and frequencies. (The forwards transitions are from being broken or not broken prior to each gust, to being broken or not broken after each gust). If indeed, upon experimentation this turn out to be so, then of course the backwards transition frequencies will not be in invariant.

But let us suppose that we do not find invariant forwards transition frequencies. Suppose that

we find that as time progresses the proportion of the remaining strings that break increases noticeably. What would we conclude? We would conclude that there is some sense in which the strings get stretched or altered by the previously applied wind gusts, so that they become structurally weaker, and hence more prone to breaking. I.e. we assume that there is a factor, previously unobserved and unknown, which does not develop chaotically, which also partially determines whether the string will break. Indeed, assuming the invariance of the distribution of the chaotically developing factors we can infer from the data, the frequencies of breaking as time develops, that there is a further factor the state of which develops in time, according to certain transition chances, upon which the forwards transition chances of the string (breaking or not breaking during a gust), depend. Indeed on the basis of the data one can conjecture a numerical scale and a time-dependence indicating the relevant state of this factor and how it develops in time in the presence of repeated wind gusts. One would then hope to find other methods for the assessment of the state of this non-chaotically developing factor, in order to confirm this conjecture further, and in order to uniquely infer the transition chances in involved.

In view of these examples let me now suggest an amendment to the account offered so far. So far I have suggested that invariant transition chances from properties A_i to properties B_j are to be expected when the further properties M_l which, in addition to A_i , determine which B_j occurs at the later time, have undergone some prior chaotic development, and hence can be expected to be distributed "uniformly with respect to future developments". But in the case of the pegboard we had a non-uniform initial distribution of the microscopic states. For if we fixed the "macroscopic" state, the small regions A and B from which we released the balls, the microscopic states compatible with these macroscopic states, the exact initial positions and velocities of the balls, were not in any clear sense uniformly distributed. More importantly, they need not even always be distributed in the same manner. For exactly how the microscopic positions around location A (or B) are distributed will generally depend on the person releasing the balls. Some people have more accurate control of their hands and fingers than others, and the distribution of the exact initial positions of the balls, even

when one is attempting to release the balls from the same spot, will differ from person to person. But I argued that the chaos in the subsequent trajectories of the balls implies that most variations in initial distributions make a negligible difference regarding the final distribution. And I now claim that peg board experiments will confirm that the forwards transition frequencies will vary only negligibly when different people, or mechanisms, perform long runs of such experiments. Thus, it seems that one does not always need a perfectly uniform and identical distribution of the "chaotically developing" factors in order to expect invariant forwards transition chances between the "non-chaotically developing" factors. If the development during the period in which the transitions take place is chaotic enough, it would seem that most initial distributions compatible with a given "non-chaotically developing factor" will imply the same, invariant, forwards transition chances between "non-chaotically developing" factors. There is a sense in which invariant forwards transition chances due to chaotic developments during the period of the transitions are more fundamental than the invariant forwards transition chances due to prior chaotic developments of the additional factors. In the former case, since the development is chaotic from each cell A_i at t_0 to the B_j at t_1 , specifying any further properties, which are not bizarre or completely fix the state (up to measure 0) at t_0 will not affect the chances of B_j at t_1 . Thus in such a case there is a clear sense in which we have said as much as we can regarding the cause of whatever property B_j it has at t_1 when we specify what property A_i it has at t_0 .

However, in the latter case, though the further factors M_i will always be uniformly distributed at the start of the transition period, the development between the A_i at t_0 and the B_j at t_1 need not be chaotic at all! Thus, it is quite possible that there are some further observable factors C_r such that the transition chances from $A_i \& C_r$ at t_0 to B_j at t_1 are not the same as the transition chances from A_i at t_0 to B_j at t_1 . Let me give an example of such a case.

Let us suppose that the producers of a particular type of metal beams have found that this type of beam appears to have a fixed and invariant chance, approximately 1 in 5, of breaking when a weight of 1000 pounds is dropped from a height of ten feet onto the middle of the beam, which is

supported at its ends. None of the controllable or observable conditions of its manufacture appear to make any difference regarding these chances. And how many previous such tests the beam has survived appears to make no difference regarding these chances. Then an observant manufacturer discovers that whether such a beam will break or not is almost entirely determined by the many very small but observable hairline cracks that occur in the beams. These cracks first appear when each beam is cooled, after it was poured into a mould, during its manufacture. And they change their shape and size each time the weight is dropped onto the beam. Eventually the manufacturers learn to predict almost perfectly whether the beam will break or not when the weight is dropped onto it, on the basis of the observable structure of the hairline cracks just prior to that drop. Unfortunately they also discover that the way in which the cracks are formed, both during the cooling and during each weight dropping test, is so chaotic that they are not able in any way to control the formation of the cracks². In such a case, the transition chance for such from being whole to being broken during each test is indeed an invariant 1 in 5. However, it is not a very fundamental transition chance, as one has a much more useful transition chance at hand for predictive purposes. And one can obviously select a subset of the beams produced at any time in such a way that the transition chances are near 0 for that subset (or near 1, or anything in between). Or, one could perhaps find some repair procedure for the cracks which will reduce the chance of breaking. If so, the transition chances would no longer be 1 in 5, but of course, it would no longer be the case that the formation of the cracks would be the result of a prior chaotic development. Thus in this case the invariance of the transition chances does not imply that one can not have a better knowledge of the chance that the beams will break, nor does it imply in any sense that we know all the causes, or even all of the observable causes, of their break-up.

4. Conclusions

One expects invariant forwards transition chances between properties A_i at t_0 and B_j at later time t_1 if the further properties M_l , which in conjunction with the A_i at t_0 determine which property

B_j obtains at t_1 , have had a prior chaotic development. In contrast one does not expect invariant backwards transition chances when the further factors N_k , which in conjunction with the B_j at t_1 determine which property A_i obtains at t_0 , develop chaotically after t_1 . One also expects invariant forwards transition chances from properties A_i at t_0 to properties B_j at t_1 if the development between times t_0 and t_1 is so chaotic that most initial chance distributions of complete states that correspond to the same property A_i , i.e. chance distributions of complete states entirely confined to cell A_i of the initial state space, develop into almost exactly the same final distribution over (incomplete, coarse) properties B_j . In contrast one does not expect invariant backwards transition chances even if the backwards development is so chaotic that most final chance distributions that correspond to property B_j , i.e. chance distributions of complete states entirely confined to cell B_j , would develop backwards in time into almost exactly the same chance distribution over (incomplete, coarse) properties A_i . The invariant forwards transition chances in the latter case are more fundamental in that no plausible further information regarding the state at time t_0 will yield any different transition chances, which is not true in the former case.

When one has non-invariant forwards transition chances from properties X_i to properties B_j , then one knows that there are further causes of properties B_j , whose frequencies vary in time. When one has found properties A_k such that the transition chances from properties A_i to properties B_j are invariant, then those A_i are causes of the B_j . If the development from each A_i to the B_j is chaotic, then those A_i are all the causes at time t_0 that we can have any hope of controlling or uncovering.

Appendix: transition chances and the inevitable
development towards equilibrium

PART I: Development towards equilibrium.

There is a large class of Markov processes with invariant forwards transition chances (stationary transition chances) which are such that any initial distribution must develop towards an equilibrium distribution as time goes to positive infinity. Let me briefly summarize some of the main results. For more detail see e.g. Doob 1953.

Ia) Discrete time, finitely many states

Definition: State j is a consequent of state i if there is some finite time interval N such that in time N there is a non-zero chance of going from i to j .

Defn: A transient state i is a state i which has a consequent j such that i is not a consequent of j .

Defn: An ergodic class of states E is a collection of non-transient states that has no consequents outside E .

Defn: Subclasses C_1, \dots, C_n are a set of cyclically moving subclasses, of cycle length n , if each state in C_j has chance 1 of next being in class C_{j+1} , where $C_{n+1} = C_1$.

Defn: A probability distribution p is an equilibrium distribution (for a set of invariant forwards transition chances) if p does not change in time (when developing according to those transition chances). Distribution p is a cyclic equilibrium distribution of cycle length n if p develops back into itself after n units of time.

Theorem: There is a unique decomposition of the set of all states into ergodic classes and a class of transient states. Each ergodic class has a unique decomposition into a set of cyclically moving subclasses.

Theorem: For each ergodic class there is a unique equilibrium distribution which is 0 outside the ergodic class (which may be cyclic). Any equilibrium distribution is a weighted mixture of these equilibrium distributions. Any initial distribution will, as time goes to positive infinity, approach such an equilibrium distribution. (The probability of the transient states will go to 0).

Ib) Discrete time, arbitrary state space

One specifies such a process by specifying invariant transition chances $p(x, A)$ from arbitrary states x to arbitrary sets of states A per unit time. This will also specify the transition chances $p^{(n)}(x, A)$ in n units of time.

Defn: A discrete time, arbitrary state space, Markov process with stationary transition chances $p(x, A)$ satisfies condition D iff there is a finite measure ϕ , an integer $n > 0$, and a non-zero ϵ , s.t. $p^{(n)}(x, A) < 1 - \epsilon$ if $\phi(A) < \epsilon$.

Basically this condition demands that no matter where you start, in the long run the probability of small sets should not remain large.

One can then define ergodic classes and transient classes in a manner analogous to the case of finitely many states. If one then applies a "reduction" procedure which at most removes a set of states of measure 0, then one can prove the following theorem of the "reduced" state space.

Thm: The state space of any discrete time Markov process with invariant transition chances which satisfies condition D can be uniquely decomposed into a set of ergodic classes and a transient class. Each ergodic class has a unique equilibrium distribution confined to that class (which may be cyclic). As time goes to infinity any initial distribution will approach a weighted mixture of these equilibrium distributions.

Ic) Continuous time, finitely many states

Each such process has a unique decomposition into ergodic classes, and a transient class. For each ergodic class there is a unique probability distribution confined to that class. Any initial distribution will approach a weighted mixture of these equilibrium distribution exponentially fast. (There can be no cyclic equilibrium distributions in this case)

Id) Continuous time, arbitrary state space

One specifies such a process by giving transition chances $p(x,A,t)$ for arbitrary states x , sets of states A , and time intervals t .

Defn: Condition D is satisfied if there is a finite measure ϕ of Borel sets of the state space, an $\epsilon > 0$, and a $\tau > 0$, such that $p(x,A,t) < 1 - \epsilon$ if $\phi(A) < \epsilon$.

Thm: There is a unique decomposition of the state space into ergodic classes and a transient class. For each ergodic class there is a unique equilibrium distribution confined to that class. Any initial distribution will approach a weighted mixture of these equilibrium distribution exponentially fast.

PART II: Invariant transition chances

In this section I want to show that invariant forwards transition chances rule out invariant backwards transition chances. To be more precise, I will show that if either one has finitely many states, or the forwards transition chances satisfy condition D, then either one has an equilibrium distribution at all times, or the backwards transition chances are not invariant.

The basic idea of this argument is very simple. We know that any probability distribution p_t governed by such forwards transition chances must develop towards a distribution p which is an equilibrium distribution for these forwards transition chances. As time goes to positive infinity p_t must get closer and closer to this equilibrium distribution p , and the rate of change of p_t will go to 0. Now suppose that p (in the backwards direction of time) develops according to invariant backwards transition chances. It will then also be true that as t goes to positive infinity the rate of change of p_t (in the backwards direction of time) according to those invariant backwards transition chances goes to 0. For the rate of change of p_t in the backwards direction of time also goes to 0 as t goes to positive infinity (near positive infinity p_t simply does not change very much, it becomes almost constant). This seems to imply that p must also be an equilibrium distribution for the invariant backwards transition chances. But that could not possibly be so. For, in the backwards direction of time, p_t develops away from this equilibrium distribution, even though a process governed by invariant backwards transition chances must develop towards such an equilibrium distribution in the backwards direction of time. Thus we have an inconsistency, and we can infer that we can not have invariant backwards transition chances unless we were in an equilibrium distribution at all times. Now let me state this argument more precisely.

IIa Discrete space, discrete or continuous time

First let me show that arbitrarily close probability distributions have arbitrarily close rates of change for discrete space Markov processes with stationary transition chances. To be more precise let me show that for any given finite time interval δt and any given $\epsilon > 0$, one can find a nonzero β such that, if the sum of the absolute value of the probability differences $\sum_n |P'(S_n,t) - P(S_n,t)| < \beta$, then $P'(S_n,t+\delta t) - P(S_n,t+\delta t) < \epsilon$.

Let $T(S_m,t+\delta t/S_n,t)$ be the transition chances over period δt , and let us abbreviate them as $T(m/n)$. (If time is discrete let δt be the unit of time in question, if time is continuous choose some finite time interval δt , and let T be the transition chances over that time period). Then $P(n,t+\delta t) = \sum_m T(n/m)P(m,t)$, and similarly $P'(n,t+\delta t) = \sum_m T(n/m)P'(m,t)$. Now let $P'(m,t) = P(m,t) + \alpha(m,t)$. Then

$P'(n,t+\delta t) = \sum_m T(n/m)(P(m,t) + \alpha(m,t)) = \sum_m T(n/m)P(m,t) + \sum_m T(n/m)\alpha(m,t) = P(n,t+\delta t) + \sum_m T(n/m)\alpha(m,t)$. Now suppose $\sum_m |\alpha(m,t)|$ goes to 0. Since all terms $T(n/m)$ are between 0 and 1 (inclusive) it follows that $\sum_m T(n/m)|\alpha(m,t)|$ will go to 0. But if $\sum_m T(n/m)|\alpha(m,t)|$ goes to 0, then so must $\sum_m T(n/m)\alpha(m,t)$. Thus for any time interval δt , one can make $P'(n,t+\delta t)$ arbitrarily close to $P(n,t+\delta t)$ by making distributions $P(n,t)$ and $P'(n,t)$ arbitrarily close (in the sense that the sum of the absolute probability differences is arbitrarily small).

Now let us argue by reductio. Suppose that we have a distribution $P_t(S_n)$ over a discrete set of states S_n , developing according to invariant forwards transition chances FTC towards a distribution $P_e(S_n)$ which is an equilibrium distribution for these forwards transition chances. That is to say, we have a 1-parameter family of probability distributions which gets arbitrarily close to

P_e as t goes to infinity (t can be a discrete or a continuous parameter). Suppose also that it develops in the backwards direction of time according to invariant backwards transition chances BTC. Suppose that P_e is not an equilibrium distribution for the backwards transition chances BTC. This means that if we let P_e develop according to these backwards transition chances for some finite time period δt then the probabilities of at least some states S_k will have changed a finite amount. But we also know that arbitrarily close distributions develop at arbitrarily close rates, so we can choose an ϵ such that if $\sum_n |P'(S_n) - P_e(S_n)| < \epsilon$, then, when we develop P' according to BTC over period δt , $P'(S_k)$ will change by some finite amount μ . On the other hand we have a 1-parameter family of probability distributions P_t which gets closer and closer to P_e such that its rate of change over a finite period δt goes to 0 as t goes to infinity. This implies that we can find a P_t such that $\sum_n |P_t(S_n) - P_e(S_n)| < \epsilon$, and $P'(S_k)$ will change by less than μ if developed according to BTC for period δt . This is inconsistent, so we have shown that one can not have invariant backwards transition chances when we have invariant forwards chances which compel a development towards an equilibrium distribution.

Ic Continuous state space, discrete or continuous time

This proof is essentially the same as the previous one. Let me start by showing that close distributions have close rates of change. Probability density $p(A, t + \delta t) = \int T(x, A, \delta t) p(x, t) dx$. Suppose that $p'(A, t) = p(A, t) + \epsilon(x, t)$. Then $p'(A, t + \delta t) - p(A, t + \delta t) = \int T(x, A, \delta t) \epsilon(x, t) dx$. Now suppose $\int |\epsilon(x, t)| dx = \mu$, for some finite μ . Since $|\epsilon(x, t)|$ is everywhere positive, $|\epsilon(x, t)|/\mu$ can be regarded as a probability density (it integrates to 1, and is positive everywhere). Thus $\int T(x, A, \delta t) |\epsilon(x, t)|/\mu dx$ is the probability that area A would have at $t + \delta t$ if we started with probability density $|\epsilon(x, t)|/\mu$. So $\int T(x, A, \delta t) |\epsilon(x, t)|/\mu dx \leq 1$. So $\int T(x, A, \delta t) \epsilon(x, t) dx \leq \mu$. Since $\int T(x, A, \delta t) \epsilon(x, t) dx \leq \int T(x, A, \delta t) |\epsilon(x, t)| dx$ it follows that $\int T(x, A, \delta t) \epsilon(x, t) dx \leq \mu$. So in this case we have also shown that arbitrarily close density distributions develop into arbitrarily close density distributions. If we now assume that the equilibrium distribution that is approached as time goes to positive infinity is not an equilibrium distribution for the backwards transition chances in the sense that the probability of some non-zero measure area A will not remain invariant, then we can derive an inconsistency as in the previous proof. Thus we have shown that in this case too one can not have invariant backwards transition chances when one has invariant forwards transition chances that lead to an equilibrium distribution.

REFERENCE

Doob, J. L. (1953): Stochastic Processes. Wiley & sons, New York.

NOTES

1 . Here I am ignoring the fact that the final position does not just depend on the initial position along 1 dimension, but depends on the original position and velocity along 3 dimensions each (for a given fixed environment). In fact, assuming that the laws of development are continuous, the later distribution would be an apparently random, but simply connected, "spiders web" with tentacles in almost every region no matter how small. However, the argument of this section would remain the same, so I will stick with the simpler picture. In the simpler picture I have made each final chop the same height (and flat on top), because I am assuming a Liouville development in the sense that volume and density of the distribution is preserved as it flows and gets carved into "chops" in time.

2 . One might object that if indeed hairline cracks are almost always formed, then the prior process in which these cracks were formed can not really have been chaotic. For the fact that such cracks, such relatively large scale structures, are present in almost all of the beams indicates that the beams are clearly in some highly non-uniform non-equilibrium state. However, in the first place it is not clear that this last claim is true. For it may very well be that for almost all microscopic configurations of beams of this type of metal, density, size, shape and temperature one has such hairline cracks. In the second place the only demand that I made for processes of type i) was that the factors determining whether the beams will break or not underwent a prior chaotic development. So, in the first place, I am only demanding a prior chaotic development of certain factors, not of everything. Secondly, although I am expecting, on the grounds of this prior chaotic development, some uniform, equilibrium distribution of these cracks, it may very well depend on what procedure is used for making the beam whether there are such cracks at all, or what their distribution is. As long as in this procedure, i.e. for this prior dynamics, the exact state of the factors which determine subsequent breaking or non-breaking are the end-result of a chaotic process, one can expect invariant transition chances during the subsequent weight-dropping test.