Entropy and the Direction of Natural Change

Binny J Cherayil

Binny J Cherayil is an Associate Professor at Department of Physical Chemistry, Indian Institute of Science, Bangalore. He works on the statistical mechanics of polymers.

Plus ça change ...

For the king of Corinth who was condemned forever to roll a large rock repeatedly up a hill only to have it roll down again when he had reached the summit, it would have been natural to conclude that the more things change the more they stay the same. But outside the pages of Greek myth, the world doesn't quite work this way. More often than not, the more things change the more they become different. A balloon filled with a coloured gas, for instance, will release its contents to the air in a room as soon as its mouth is opened or its skin punctured, but an empty balloon in the same room with just the same amount of coloured gas will never fill up with the gas on its own. Irreversibility is generally the norm, not the exception, to the way ordinary matter behaves, and there are certain directions of change that we know instinctively are allowed and certain others that we know are not. Why should that be? After all, transformations between different kinds of matter ultimately involve nothing more than the rearrangement of atoms and molecules, and that requires just the same expenditure of work in one direction as it does in another. We know this because the total supply of energy in the Universe is fixed and unchanging, so if some of it is used up at point A, exactly the same amount will be returned at point B. Neverthless, emptying a balloon is easy, filling it up is hard. Mere respect for the constancy of energy is apparently not enough to prevent nature's preference for certain directions of change over others. What then is?

Irreversibility is generally the norm, not the exception, to the way ordinary matter behaves, and there are certain directions of change that we know instinctively are allowed and certain others that we know are not.

The answer lies in the properties of a quantity that though well known is often not well understood. We shall now try to learn what that quantity is and what it tells us about the direction of natural change.

Definitions

If we're interested not so much in deflating balloons as in any kind of physical process, we need a more general set of terms to describe the changes it may undergo. It's convenient therefore to refer to any object whose behavior we're interested in as a system and everything else as its surroundings, the two being separated by a boundary, which we'll understand to mean a barrier that limits the degree of contact between them.

System and surroundings generally have physical attributes that we can perceive or that can be made perceptible by certain methods of measurement. A list of these attributes (or properties) tells us what state the system or surroundings is in. It's often a matter of judgement how big this list has to be before we know enough to identify the state unambiguously.

System and surroundings generally have physical attributes that we can perceive or that can be made perceptible by certain methods of measurement. A list of these attributes (or *properties*) tells us what *state* the system or surroundings is in.

In the example of the inflated balloon, we can think of the coloured gas (bromine, say) as the system, the skin of the balloon as the boundary and the room as the surroundings. System and surroundings both have definite physical attributes that when catalogued provide a precise description of each. For the system, it would be enough to say that it was made up of the gas bromine, that the gas was at a certain temperature and pressure, and that it weighed so many grams. That limited information would be enough for someone else in some other part of the world to faithfully recreate the conditions that define the kinds of changes we're seeking to monitor. Moreover, by making the room big enough, and by sealing it off from all the other rooms in the house, we can essentially guarantee that everything that happens before and after we have punctured the balloon happens wholly within the confines of the room, and nowhere else. In effect, system and surroundings in this instance constitute a composite system that is closed with respect to the rest of the universe. The composite system is in a certain state X before the skin of the balloon is opened, and it's in a certain state Y thereafter.

Expressed in these terms, the problem of why a filled balloon spontaneously deflates is really the problem of why a closed composite system is driven to a definite new equilibrium state when an internal constraint is removed.

Microscopic vs. Macroscopic

But we're no closer to solving the problem. It turns out that to make further progress, we need to say more about the system at the microscopic level. It's not enough to describe the system solely in terms of properties like temperature and pressure and volume, even if that description happens to be complete at the macroscopic level. But what constitutes a microscopic description of the system? The technical answer is its phase space, but it's simpler to think about this question in terms of analogies.

Consider, for example, a set of 3 dice. Each die has between 1 and 6 pips on its faces. When the dice are rolled together, the system of 3 dice can be said to produce a definite microstate in which the first die shows, say, a 5, the second shows a 3 and the third shows a 6. There are clearly a total of $6 \times 6 \times 6$ such microstates. Some of those microstates may be of greater significance than others in certain games of chance. If the dice showed a total count of, say 10, for example, a player might win a certain sum of money. There are obviously far fewer of these special throws than the 216 combinations that are possible. The special throws (where the total number of pips is predetermined) can be thought of as macrostates of the system; they represent a bulk property of the system in that they are unconcerned with which die has what number of pips on its face so long as the total works out

The number of microstates that correspond to or are consistent with a given macrostate is usually called the degeneracy or the multiplicity of the macrostate.

to be some definite number. This means, of course, that while there is just the single macrostate corresponding to the total 10, there are quite a few microstates that could produce it. 6, 3, 1 is one; 4, 4, 2 is another; and 2, 5, 3 is a third. The number of microstates that correspond to or are consistent with a given macrostate is usually called the degeneracy or the multiplicity of the macrostate. Table 1 lists the microstates that produce the macrostates 3, 4 and 5, along with the corresponding degeneracies.

It's no surprise that these macrostates don't all have the same degeneracy. That means that any given throw of the set of 3 dice is more likely to lead to certain macrostates than they are to others. This notion of likelihood can be made more precise by defining the probability of producing a given macrostate, p(x), as the ratio

Table 1. The states of a system of 3 dice. Microstates are shown in the first column; they are the detailed dice configurations showing the number of pips on each die. Macrostates are shown in the second column; they correspond to the total of the three dice in the first column. The degeneracy of a given macrostate is shown in the third column. The probability of throwing a given total is shown in the fourth column.

Microstate	Macrostate	Degeneracy	Probability
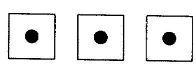	3	1	1/216
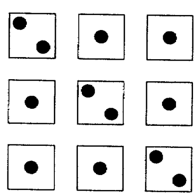	4	3	3/216=1/72
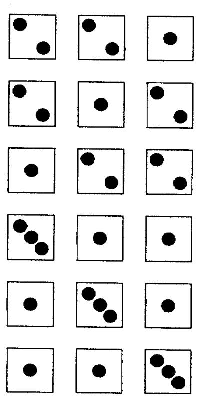	5	6	6/216=1/36

of the degeneracy of that state – $\Omega(x)$ to the total number of possible microstates – $\Omega(\text{total})$. Thus, for instance,

$$p(10) = \frac{\Omega(10)}{\Omega(\text{total})} = \frac{27}{216} = \frac{1}{8}. \qquad (1)$$

Similarly, $p(14) = 5/72$, and so on for any other macrostate of the system. Table 1 also lists the probabilities of observing the macrostates 3, 4 and 5.

Interacting Systems

What has all this to do with natural processes? At the moment, not much, but the dice are meant to represent the particles that make up a real chunk of matter (such as bromine gas, for example), the 6 numbers on their faces are meant to represent the internal energy states that these particles can be in, and the total of the three numbers is meant to represent the energy that a bulk sample of the gas might have. The analogy is actually not all that far-fetched, but it's by no means perfect. Any bulk sample of real matter will obviously have far more than 3 constituent particles (it will typically have on the order of 10^{23}), and these particles will have much more than 6 internal states (infinitely more, for all practical purposes because of the different electronic, vibrational and rotational quantum numbers each particle may assume.) Moreover, it's seldom the case that the state of the system as a whole can be specified merely by specifying the internal energy state of each particle. That can only be done when the particles are independent. When they're not, the microstates have to be specified differently, but the differences are only a matter of detail, and they don't really affect the argument or its conclusions.

So we'll continue to pretend that 3 dice with 6 independent internal states are a satisfactory representation of a bulk sample of matter, in this case bromine. Along the same lines, and after further suspension of disbelief, we'll represent the surroundings (the air in the room) by a set of 3 more dice. We'll call these two systems B and R (for balloon and room). The macrostates of both B and R range from 3 to 18 (there being 6^3 microstates that give rise to them). If we combined the two sets of dice to form a larger system containing 6 dice, the corresponding macrostates would then lie between 6 and 36, and would be produced by $6^6 = 6^3 \times 6^3$ microstates. It's convenient to think of B and R as a composite system because it's a simple matter to bring together the two containers that held them separately before, just as the filled balloon and the room form a composite system because the two can be brought together by opening the mouth of the balloon. But remember that the composite system of balloon and room was closed; if the composite system formed by B and R is to be similarly closed, the total face count of the dice before and after they are mixed must be the same. For definiteness, therefore, let's suppose that this total is 21, and that before mixing, B is in the macrostate 3 and R is in the macrostate 18. This choice of parameters is meant to reflect the disparity in energy between the contents of the inflated balloon and those of the room, the room accounting for most of the available energy.

It doesn't matter at this point just how the dice found themselves in these two macrostates; it's enough that things can be so arranged. Of greater interest is what happens thereafter, when the two sets of dice are combined and then rolled together. Most throws will not produce a total of 21, but whenever any does, we record the total from the three dice that came from B, which we may suppose have previously been coloured blue to distinguish them from the R dice, which have similarly been coloured red.

As we roll the dice, and jot down totals, we will discover that certain macrostates of B come up more often than others. That's clearly because the degeneracies of such

Table 2. The states of a system of 3 blue (B) dice in a composite system of 3 blue and 3 red (R) dice in which the total of the 6 dice is always 21. The first column shows the macrostates of B, the second their degeneracies, and the third their probabilities.

Macrostate	Degeneracy	Probability
3	1	2.31×10^{-4}
4	9	2.08×10^{-3}
5	36	8.31×10^{-3}
6	100	0.023
7	225	0.052
8	441	0.10
9	625	0.14
10	729	0.17
11	729	0.17
12	625	0.14
13	441	0.10
14	225	0.052
15	100	0.023
16	36	8.31×10^{-3}
17	9	2.08×10^{-3}
18	1	2.31×10^{-4}

states are larger. It's not difficult to determine these degeneracies beforehand, and they are shown in Table 2. From the table it's a simple matter to calculate the likelihood of the appearance of certain states. As before, it is simply the ratio of favourable outcomes to the total of all possible outcomes. These probabilities are also listed in Table 2. What is immediately apparent is that having started out in the macrostate with a count of 3, B is quite unlikely to stay there after mixing with the room (red) dice. There are so many more states that correspond to 'not-3' that the odds of one of those states showing up are so much higher. The state that most likely appears in a throw is the one with the greatest degeneracy, which in the present example is the state in which the B dice have a total of 10 and the R dice a total of 11, or vice versa. These are the states that not only have the greatest degeneracy, they are also the ones that in some sense exhibit the greatest degree of 'sharing' of the initial 'energies' in B and R; R was initially energy-rich and B energy-poor, but when the two interact, they tend to share that energy more or less equally. These observations are summarized in the graph shown in Figure 1, which illustrates how the initial states of B and R, over the course of many throws, will tend to be seen most often around the state of greatest degeneracy.

The graph in Figure 1 shows a fairly broad maximum, which tells us that there are actually quite a few macrostates in the vicinity of the true maximum that have quite large degeneracies themselves. How would this picture change if the game were played not with 6 dice but with, say, ten dice, five red and five blue? To find out, we can imagine, as before, that the system starts out in a state where the blue and red dice have fairly disparate totals. Let's assume, therefore, that the 5 B dice each shows a 1, and the 5 R dice each shows a 6, for a grand total of 35 when the dice are mixed and all the pips are added together. Since the system is closed, the totals of subsequent throws of the dice are recorded only if they happen to be 35. After many throws, it will become quite apparent that only very rarely (and may be not at all if we quickly tire of the game) will we ever see the B and R dice in the all-1 and all-6 states they were in at the beginning. Table 3, which shows the number of microstates that correspond to a given B total when the overall count is 35, tells us why. There are just so many more microstates that are not all-1 that their chances of being observed are so much greater, exactly the situation we found for the system of 3 red and 3 blue dice when the total face count was 21. But now there are even more microstates that are distinct from the initial configuration that it's even more unlikely than was the case with 6 dice that the unique starting state will ever reappear during the course of the game. The microstates that are likely to be seen are those that have comparable totals on the B and R dice. The implication, again, is that the initial state tends to evolve towards a condition of maximal sharing of the available 'energy' of the two sets of dice.

We could continue in this fashion to look at other still

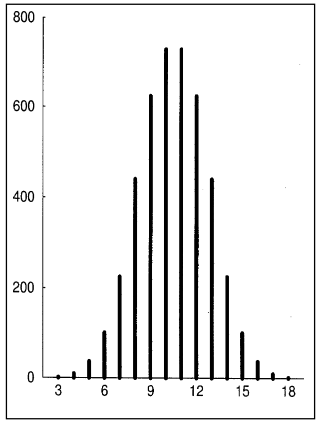

Figure 1. A graphical depiction of the data in Table 2. The numbers in the first column of Table 2 are shown on the x-axis (as the 'energy' of B), and the numbers in the second column are shown on the y-axis.

The initial state tends to evolve towards a condition of maximal sharing of the available 'energy' of the two sets of dice.

Table 3. The states of a system of 5 blue dice in a composite system of 5 blue and 5 red dice in which the total is 35. The three columns in this table have the same meaning as the corresponding columns in Table 2.

Macrostate	Degeneracy	Probability
5	1	2.28×10^{-7}
6	25	5.69×10^{-6}
7	225	5.12×10^{-5}
8	1225	2.79×10^{-4}
9	4900	1.11×10^{-3}
10	15876	3.61×10^{-3}
11	42025	9.56×10^{-3}
12	93025	0.0212
13	176400	0.0401
14	291600	0.0663
15	423801	0.0964
16	540225	0.123
17	608400	0.138
18	608400	0.138
19	540225	0.123
20	423801	0.0964
21	291600	0.0663
22	176400	0.0401
23	93025	0.0212
24	42025	9.56×10^{-3}
25	15876	3.61×10^{-3}
26	4900	1.11×10^{-3}
27	1225	2.79×10^{-4}
28	225	5.12×10^{-5}
29	25	5.69×10^{-6}
30	1	2.28×10^{-7}

larger systems, (systems that had on the order of 10^{23} dice, for instance), and in all these cases we would find that the new equilibrium state that a composite system tends to evolve to in the absence of an internal constraint is the one with the largest degeneracy and the greatest degree of spreading and sharing of the energy. Indeed, with 10^{23} dice, the graph of energy versus degeneracy analogous to Figure 1 would be a single sharp spike with next to no thickness at all, reflecting the near-certainty of any initial state eventually ending up – under conditions of constant total energy – in one of the cluster of states with the largest degeneracies.

A Second Law of Nature

The above remarks are certainly true of the dice in our imaginary game of chance, but there is no reason why they shouldn't be true of gas particles, for instance, or of molecules in a liquid, or of any of the other constituent elements of the matter around us. These elements may be much smaller (by far) than dice, and they may have more complicated internal states, and they may not all be independent of each other, but as we've said before, these are only points of detail: the basic argument stays the same. So we can almost certainly assert that any closed composite system that changes its state after the removal of an internal constraint ends up in the state of greatest degeneracy. (We may not know how to compute this degeneracy, but that doesn't really matter.) Since the Universe represents the ultimate closed system (by definition there's nothing beyond it), we can enunciate this principle somewhat more pithily as the multiplicity of the Universe increases. Multiplicities in general are huge numbers, however, and in actual calculations, it would be more convenient if the numbers were smaller; one way to ensure this is to use the logarithm of the multiplicity as the measure of the extent of degeneracy.[1] This quantity, when multiplied by a suitable proportionality constant k, has a special name: entropy. Mathematically, the definition of entropy reads

So we can almost certainly assert that any closed composite system that changes its state after the removal of an internal constraint ends up in the state of greatest degeneracy.

$$S = k \ln - ; \qquad (2)$$

where k is the proportionality constant that gives S the units that make it 'compatible' with certain other properties of the system, like its energy and temperature. In terms of the entropy, then, our putative law of nature is the statement that

The entropy of the Universe increases

So here finally is the explanation of nature's preference for certain directions of change: she selects those directions that produce an overall increase in the entropy.

[1] Convenience isn't the sole criterion – in fact, it really isn't a criterion at all – for introducing the logarithm at this stage; the real reason has to do with the mathematical property that the log of the product of two numbers is the sum of the logs of those numbers.

The filled balloon that spontaneously deflates does so because there are so many more states in which the molecules of bromine are interspersed more or less uniformly amongst the molecules of air than there are states in which the bromine is confined to a small part of the room. These latter states aren't forbidden; they are simply overwhelmingly improbable.

Entropy is often said to be a measure of disorder, which is a fairly useful way of thinking about this quantity, but it can be misleading, because disorder is often understood to refer to *spatial* disorder.

Entropy is often said to be a measure of disorder, which is a fairly useful way of thinking about this quantity, but it can be misleading, because disorder is often understood to refer to spatial disorder. Consider, for example, two equal volumes of the same gas at two different temperatures. Which would you say has the greater entropy? The idea of entropy as disorder doesn't really help here since in both samples of gas the molecules seem to rush around equally chaotically, with neither exhibiting any obviously greater degree of spatial randomness. But the interpretation of entropy as the extent of degeneracy tells us at once that the gas at higher temperature has the greater entropy because it has more ways than the other gas of distributing the thermal energy available to it amongst the internal energy states of its molecules. On the basis of the disorder interpretation, one might also be tempted to conclude that a beaker of crushed ice has more entropy than a beaker of the same amount of liquid water, but of course one would be wrong. The greater disorder of the crushed ice is only apparent; it is the far greater degeneracy – 'freedom', loosely speaking, rather than disorder – of the liquid that is the crucial element in the analysis. On the other hand, if nothing is really immune from the operation of the law of entropy increase, as we now believe, why would ice ever form in the first place, crushed or otherwise? The reason is that it is the entropy change of the system and surroundings that must always increase, not that of just one or the other. If entropy decreases at P, it has to increase at Q, and by a sufficient amount that the net change in

The reason is that it is the entropy change of the system *and* surroundings that must always increase, not that of just one or the other.

the entropy is positive. (Or zero, for perfectly reversible processes.) Keeping track of the entropy changes that occur in both system and surroundings can sometimes be difficult, (especially if the surroundings encompass just about everything in the Universe), so it would be far more convenient if there were a way to reach the same conclusions about the direction of natural change from an examination of the system alone. As it turns out, there is, but that is another story...

Notes

Much of the material in this article has been taken from Thermodynamics and an Introduction to Thermostatistics by H B Callen and from An Introduction to Thermal Physics by D Schroeder.

Acknowledgements

I am grateful to Ram Seshadri, IISc, for a critical reading of the manuscript, and for constructing Table 2 and Figure 1.

Address for Correspondence
Binny J Cherayil
Department of Inorganic and Physical Chemistry
Indian Institute of Science
Bangalore 560012, India.

"One quantum notion that mystifies the novice is the 'wave-particle duality.' Does light consist of a beam particles or is it a wave phenomenon? the question is hundreds of years old. Newton thought light was probably a stream of particles. Maxwell seemed to answer the problem decisively by showing light to be an electromagnetic wave. Yet Einstein in 1905, demonstrated that under some circumstances light behaves as if it were a beam of discrete particles, which are now called photons."

Sheldon L Glashow
Higgins Professor of Physics at Harvard University
Nobel Laureate, 1979