Pergamon

Stud. Hist. Phil. Mod. Phys., Vol. 32, No. 4, pp. 581–594, 2001
© 2001 Elsevier Science Ltd. All rights reserved.
Printed in Great Britain
1355-2198/01/$ - see front matter

Ergodic Theory, Interpretations of Probability and the Foundations of Statistical Mechanics

*Janneke van Lith**

The traditional use of ergodic theory in the foundations of equilibrium statistical mechanics is that it provides a link between thermodynamic observables and microcanonical probabilities. First of all, the ergodic theorem demonstrates the equality of microcanonical phase averages and infinite time averages (albeit for a special class of systems, and up to a measure zero set of exceptions). Secondly, one argues that actual measurements of thermodynamic quantities yield time averaged quantities, since measurements take a long time. The combination of these two points is held to be an explanation why calculating microcanonical phase averages is a successful algorithm for predicting the values of thermodynamic observables. It is also well known that this account is problematic.

This survey intends to show that ergodic theory nevertheless may have important roles to play, and it explores three other uses of ergodic theory. Particular attention is paid, firstly, to the relevance of specific interpretations of probability, and secondly, to the way in which the concern with systems in thermal equilibrium is translated into probabilistic language. With respect to the latter point, it is argued that equilibrium should not be represented as a stationary probability distribution as is standardly done; instead, a weaker definition is presented. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Ergodic Theory; Statistical Mechanics; Stationarity; Equilibrium; Interpretation of Probability.

*Institute for History and Foundations of Science, P.O. Box 80.000, 3508 TA Utrecht, The Netherlands (*e-mail:* j.h.vanlith@phys.uu.nl).

PII: S1355-2198(01)00027-2

1. Introduction

This paper investigates the role that ergodic theory can play in the foundations of equilibrium statistical mechanics. Historically, the mathematical theory of ergodic theory developed in the early twentieth century in close connection with foundational issues in statistical mechanics. Its original role was to establish a connection between ensemble functions (phase averages) and properties of individual systems (time averages). Whether it gives a valid justification for the use of ensembles in statistical physics is however much disputed.

In this paper I will discuss three ergodic approaches that differ from the standard ergodic approach. Two of them are meant as support of specific interpretations of probability in statistical mechanics, namely of the time average interpretation, and the personalist interpretation, respectively. Both are advocated by Von Plato and Guttmann (Von Plato, 1988, 1989; Guttmann, 1999). The third approach, originally put forward by Malament and Zabell and elaborated by Vranas, is aimed at the same goal as the standard ergodic approach, namely to explain the success of the phase averaging method, but uses a different line of argument (Malament and Zabell, 1980; Vranas, 1998). All three approaches concern the foundations of equilibrium theory, and thus have nothing to do with the issue of irreversibility, which is yet another area where ergodic theory may have a role.

A second aim of this paper is to investigate the relevance of the interpretation of probability to the above-mentioned ergodic approaches. It is natural to associate ergodic theory with objective interpretations of probability. This is because with the use of ergodic theory a connection can be established between probability measures on the one hand, and objective features of real world systems on the other. However, ergodic theory can be useful also with other interpretations of probability, as is shown by the use of the ergodic decomposition theorem as support for the personalist interpretation of probability.

The interpretation of probability and the ergodic approaches are clearly connected in the sense that ergodic theory has been invoked as putative support for two distinct interpretations of probability: the time average interpretation, and the personalist interpretation. Another sense in which the interpretation of probability is relevant is, I will argue, that the plausibility of certain assumptions that are made in the mentioned ergodic approaches depends on the interpretation of probability.

In order to illustrate both the diversity of foundational roles of ergodic theory and the importance of the interpretation of probability, I will single out one particular assumption about the probability distribution that plays a role in all three ergodic approaches, namely stationarity. A probability distribution is stationary if it is constant at all fixed points in phase space. This is usually taken to reflect the fact that the system described by the probability distribution is in equilibrium. However, as I will argue, stationarity of the ensemble is not the right way to account for the system being in equilibrium. I will present another, weaker condition on the probability distribution as representing thermal equilibrium, and will investigate what effect this modification has on the three ergodic approaches.

This paper is structured as follows. In Section 2, I will review the traditional role of ergodic theory in the foundations of statistical mechanics and the reasons why it is problematic. In Section 3, I will discuss the three alternative uses of ergodic theory in the foundations of equilibrium statistical mechanics. In Section 4, I will highlight the notion of stationarity, and argue for a weakened account of equilibrium.

2. The Standard Ergodic Approach

2.1. Standard role of ergodic theory in the foundations of statistical mechanics

The traditional use of ergodic theory in the foundations of equilibrium statistical mechanics is to make a connection between the ensembles used in statistical mechanics and properties of single systems. More specifically, ergodic theory is invoked to solve the *ergodic problem*, which is to demonstrate the equality of infinite time averages and phase averages, i.e. expectation values with respect to the microcanonical measure on phase space. A related goal is to explain the success of microcanonical phase averages, for which a solution of the ergodic problem would form the first step. The second step is then to argue that time averages are equal to the results of a macroscopic measurement. This is usually done by pointing to the fact that measurements take an amount of time which is long compared to microscopic relaxation times. The argument is then that a single measurement yields an average of the phase function over this time. These two steps taken together imply that the results of macroscopic measurements are equal to microcanonical phase averages.

Originally, the ergodic problem was attacked by means of the so-called ergodic hypothesis, which states that a system that is left to itself will pass through all the phase points compatible with its total energy. Then, since a point in phase space cannot lie on more than one phase trajectory, all systems with the same value of the total energy will follow the same path, which fills the phase space completely. At a specific moment in time different systems may be in different points on this path, but averages over infinite times are equal. Thus, if the ergodic hypothesis is satisfied, time averages will be the same for all systems with the same total energy, and this solves the ergodic problem. Unfortunately it is generally impossible to satisfy the hypothesis. One-to-one and continuous transformations leave the number of dimensions invariant, from which it follows that a one-dimensional curve cannot fill the whole energy surface, except in the trivial case when the latter is one-dimensional itself.

Although a single phase trajectory cannot visit every point on the energy surface, it might come arbitrarily close to every point. For this behaviour the term 'quasi-ergodicity' was coined by the Ehrenfests (Ehrenfest and Ehrenfest-Afanassjewa, 1912). However, it has never been proved that quasi-ergodicity is sufficient to solve the ergodic problem. Modern ergodic theory has developed along a different route, which is framed in a measure-theoretical setting.

Its objects of study are the so-called measure preserving dynamical systems (m.p.d.s.) $\langle \Gamma, \mathscr{B}, \mu, T_t \rangle$. In the applications to statistical mechanics the elements of a m.p.d.s. are identified as follows. Γ is the phase space, or the accessible part of it which is usually an energy surface. A point $x = (\vec{p}^N, \vec{q}^N)$ in Γ represents the positions and momenta of all the N particles in the system. $\mathscr{B}$ is the standard Borel σ-algebra. μ is the normalised Lebesgue-measure restricted to Γ; if Γ is an energy hypersurface, this is the microcanonical measure. Finally, T_t is the Hamiltonian evolution, parametrised by the time t. It follows from Liouville's theorem that these choices indeed lead to a dynamical system which preserves measure.

The answer to the ergodic problem as given by modern ergodic theory is contained in Birkhoff's ergodic theorem, and one of its corollaries. The ergodic theorem demonstrates the existence of infinite time averages for almost all initial conditions; the corollary demonstrates that the values of infinite time averages are equal to phase averages in the standard measure, if a dynamical condition called 'metrical transitivity' (also simply called 'ergodicity') is satisfied, again for almost all initial conditions.

The next step in the argument is the claim that measurement results yield infinite time averages of phase functions, because measurements take a long time compared to the relevant microscopic relaxation times. Then, combining this with the corollary of Birkhoff's ergodic theorem, the standard ergodic approach demonstrates that for metrically transitive systems, measurement results are almost always equal to microcanonical averages. This is the explanation it provides for the success of the averaging method.

2.2. *Problems with the standard ergodic approach*

Several problems surround the standard ergodic approach (see for example Earman and Rédei, 1996; Sklar, 1973, 1993; Jaynes, 1967). Below I will mention four of them. The last one (connected to the way equilibrium is represented) I believe has not been raised before.

The measure zero problem

The ergodic theorem and its corollary hold for almost all points in the measure theoretic sense of the phrase 'almost all'. It is tempting to conclude that since exceptions to those statements get probability measure zero, they can be neglected in practice. However, this does not follow straightforwardly but should be argued for. Especially so, since the ergodic theorems are meant to *demonstrate* that the microcanonical measure is the appropriate probability measure in a statistical mechanical treatment of isolated systems in equilibrium. The statement that having microcanonical measure zero implies being negligible in practice is thus part of the goals that the ergodic approach tries to achieve.

There have been several attempts to solve the measure zero problem (for a discussion see Sklar, 1993, pp. 182–188). One other interesting, though unsuccessful, approach is presented in Guttmann (1999), who has tried to demonstrate the exceptionless equality of phase and time averages by showing that both phase and time measures are equal to the so-called Haar measure. However, this approach fails because of its limited applicability; as Guttmann himself notes, it works only for a two-dimensional phase space, which is not very surprising because then the good old ergodic hypothesis can be satisfied. A more promising, but still not unproblematic, approach to solve the measure zero problem is presented in Malament and Zabell (1980), where absolute continuity[1] is shown to follow from another property, called translation continuity. The latter property is, Malament and Zabell claim, easier to justify.

The restriction to ergodic systems

It seems typically to be the case that the systems in classical statistical mechanics are not metrically transitive (Earman and Rédei, 1996; Vranas, 1998; Wightman, 1985). The important thing to notice is that statistical mechanics 'works' for some systems that are not metrically transitive as well (one may think of systems with small KAM-tori). Thus, the account by means of the ergodic theorems is by no means sufficient to capture all physical systems for which an explanation of the success of the phase averaging method is needed. But why should we not be happy with a partial result, that is, with an explanation of the success of the phase averaging method for the special class of metrically transitive systems? The answer, given by Earman and Rédei, is convincing. If there is an explanation for systems that are not metrically transitive, it seems reasonable to expect it to be a good explanation also for systems that are metrically transitive. But then the ergodic approach will be by-passed completely; after all, the equality of phase and time averages holds if *and only if* the system is metrically transitive.

Infinite time averages

Much of the criticism of the traditional ergodic approach is directed at the statement that measurements of macroscopic quantities yield infinite time averages of the corresponding phase function (see for instance Sklar (1973), Malament and Zabell (1980) and Jaynes (1967, p. 94)). The argument that is usually given for this is that measurements take a period of time, which is long compared to the relevant microscopic time scales (relaxation times). Thus, during a measurement the system passes through many microscopic states, and therefore the observed value will be equal to the time average.

[1] A measure ν is said to be absolutely continuous (w.r.t. Lebesgue measure μ) iff $\nu(A) = 0$ implies $\mu(A) = 0$ for all measurable sets A. Any argument to the effect that the 'actual' probability distribution is absolutely continuous would thus answer the measure zero problem.

Several arguments can be given against this line of thought. First and foremost, it is obvious that not all measurement results may be equated with infinite time averages, because it is after all possible to observe changes. Indeed, consecutive measurements on a system will in general yield different results; otherwise it would be impossible to observe any non-equilibrium process. The infinite time averages are, however, the same since the measurements are performed on a single system. Thus the claim that measurements always yield infinite time averages is simply not correct. Also, even if actual measurements take some period of time which is long compared to microscopic time scales, why would it follow that what we measure are averages over time? Finally, the value of the average over any finite period of time may differ appreciately from the value of the infinite time average. Thus, as long as one does not have any information about the rate of convergence, one cannot infer the value of the infinite time average from any measurement that takes a finite period of time. This list of objections clearly demonstrates that measurement results cannot generally be equated with infinite time averages.

Equilibrium

The standard ergodic approach, as I have presented it here, aims to explain why microcanonical averages are good predictors for observed thermodynamic quantities. It should be noted, however, that they are only good predictors for systems which are in thermal equilibrium. This implies that any explanation should take into account that systems in thermal equilibrium are being discussed; otherwise, an explanation of the success of the averaging method would be given also for cases in which this method actually is not successful. But here, surprisingly, the standard ergodic approach falls short. Nowhere in the explanation scheme is it assumed explicitly that equilibrium systems are concerned.

3. A Plurality of Ergodic Approaches

In this section I will discuss three different roles that ergodic theory can play in the foundations of equilibrium statistical mechanics. The first two are as putative support for particular interpretations of probabilities, namely the time average interpretation and the personalist interpretation, respectively. The third is aimed at exactly the same goal as the traditional ergodic approach, namely to provide an explanation of the success of the microcanonical phase averaging method; it differs in the sense that infinite time averages now play no role at all.

3.1. *Time averages*

The time average interpretation of probability is in many respects similar to the frequency interpretation. However, repetitions occur in a single system in the course of time, rather than across an ensemble of identically prepared systems. The distinctive feature is thus that the repetitions are determined by a deterministic process. A proponent of the time average interpretation of probabilities in statistical mechanics is Von Plato (1988, 1989); see also Guttmann (1999).

In the time average interpretation the probability to find a system in a certain set in phase space is by definition equal to the infinite time average of the indicator function of that set:

$$P_{x_0}(A) = \lim_{T \to \infty} \frac{1}{T} \int_0^\infty \mathbf{I}_A(T_t x_0) dt. \tag{1}$$

Thus, the probability of the set A is equal to the fraction of time that the systems spends in that region, also called the sojourn time. Note that the probability function is labelled by the initial state of the system, x_0. In general different initial states lead to different paths in phase space, and therefore also the sojourn times may depend on x_0.

There are several problems with the time average interpretation. First, the fact that repetitions are determined by a deterministic process puts pressure on the condition that the repetitions should be independent. Secondly, infinite time averages need not even exist! It may well be that the limit in (1) does not exist. Thirdly, as noted the probability of a set A depends on the initial state x_0, which is an awkward feature. Fourth, there is no obvious way to extend the application of this notion of probability to time-dependent phenomena, and thus to the more general theory of non-equilibrium statistical mechanics.

According to Von Plato ergodic theory points to cases where (some of) the mentioned problems can be overcome and thus this particular interpretation can be applied:

> the notion of ergodicity gives us a *characterization* of cases where probability as time average, and frequentist probability more generally, can be applied (Von Plato, 1988; italics in the original).

He concludes that frequentist probability can only be applied to ergodic systems. Thus, the role ergodic theory now plays in the foundations of statistical mechanics is to provide support for this particular interpretation of probability. Let us look at the four mentioned problems.

Infinite time averages, if they exist, obey the axioms of Kolmogorov. But do they also satisfy the demands of frequentist probability? Especially the condition of independent repetitions is very difficult to satisfy. In fact Von Mises, the founding father of the frequency interpretation of probability, is very clear that the time evolution of a single system does not build a *Kollektiv*, because one of the axioms of his theory of probability, the condition of random place selection (*Regellosigkeit*) is not fulfilled (Von Mises, 1931, p. 519). But whether the sampling is 'unbiased', or whether the trajectory can be seen as a sequence of independent repetitions of a random event, depends on the dynamics of the system, and here ergodic theory may be of help. If the dynamics is metrically transitive, we have 'asymptotically representative sampling' (in Von Plato's words). But only at the top of the ergodic hierarchy, for Bernoulli systems, do we have independent repetitions.

The second problem is that time averages need not exist. But Birkhoff's ergodic theorem demonstrates the μ-almost everywhere existence of infinite time averages. Thus, the existence of the probabilities as defined above is ensured for almost all starting points x_0, where 'almost all' is measured in the Liouville measure.

The third problem is that time averages generally depend on the initial state. The first corollary of the ergodic theorem shows, however, that for metrically transitive systems, time average probabilities are equal to the microcanonical measure (again with the proviso of exceptions of μ-measure zero). This means that in this case infinite time averages are independent of the initial state x_0.

The fourth problem is that the time average interpretation cannot be generalised to time-dependent phenomena. Now Von Plato is very clear that one need not pursue a single interpretation of probability in all applications of probability theory, and I agree with him. But still it would be a strange state of affairs to be compelled to use different interpretations in the single context of statistical mechanics. Indeed, how could one make sense of the statistical mechanical description of non-equilibrium phenomena, for example the approach to equilibrium?

With respect to this fourth problem ergodic theory can offer no help. In my view the impossibility of incorporating probabilities as time averages into the general non-equilibrium theory renders this interpretation untenable.

3.2. Ergodic decomposition

Another role for ergodic theory is as support for another specific interpretation of probabilities, namely personalist probabilities. According to this interpretation, with De Finetti as its main proponent, probabilities represent personal degrees of belief. An important theorem in support of this interpretation is De Finetti's representation theorem, which holds for exchangeable probabilities. Exchangeability means that $P(E_1, \ldots, E_n) = P(E_{i_1}, \ldots, E_{i_n})$, i.e. probabilities are invariant under permutations; for instance, a probability ascription in a coin tossing experiment that depends on the number of heads and tails in a sequence, but not on their order is exchangeable. De Finetti shows that exchangeable probabilities can be written as follows:

Theorem 1. (**De Finetti's representation theorem**). *$P(E_1, \ldots, E_n)$ is exchangeable for all n iff a probability density $\phi(\lambda)$ exists for which*

$$P(E_1, \ldots, E_n) = \int_0^1 \lambda^k (1-\lambda)^{n-k} \phi(\lambda) d\lambda. \tag{2}$$

Here k is the number of 'heads' and $n-k$ the number of 'tails' in the sequence of outcomes $E_1, \ldots, E_n$.

Note that the expression on the right hand side would also be arrived at in case of an objective probability λ for a coin to land 'heads', independent repetitions of coin tosses, and a prior probability distribution $\phi(\lambda)$. The importance of the theorem lies in the fact that it demonstrates that the same expression can now be obtained by assuming that one's personal degrees of belief are exchangeable. That is, one can act *as if* objective probabilities existed, and use all results from probability theory that apply to sequences of independent classically distributed trials without being committed to a belief in objective, unknown probabilities.

As pointed out by Von Plato, the ergodic decomposition theorem can do in dynamical systems theory exactly what De Finetti's representation theorem does in the case of repeated chance experiments (Von Plato, 1988).

> **Theorem 2. (Ergodic decomposition)**. *Let M be the set of all stationary measures on $\langle \Gamma, \mathscr{B}, T_t \rangle$, where $\mathscr{B}$ is bounded. Then M is non-empty and convex, and its extreme elements are ergodic.*

It follows from the ergodic decomposition theorem that any stationary measure can be written as a weighted average of ergodic measures. Thus, with the use of the first corollary of Birkhoff's ergodic theorem, it follows that they are weighted averages of time averages, and thus of objective features of individual systems. This proves that subjectivists of the De Finetti type, as long as their personal probability ascription is stationary, can do without objective probabilities; they can take the view that the latter are 'artefacts' that can be eliminated.

3.3. A uniqueness theorem

The fourth role of ergodic theory in the foundations of statistical mechanics, like the standard ergodic approach, has as its objective to explain why microcanonical averages 'work' in the sense that they successfully predict the values of thermodynamic observations (Malament and Zabell, 1980; Vranas, 1998). However, the way in which ergodic theory is used differs from the standard approach, the most important difference being that the outcomes of measurements are equated with values of phase functions, not with their infinite time averages. Also, a way to solve the measure zero problem is proposed.

In Malament and Zabell's strategy two approaches come together: on the one hand certain limit theorems which make use of the large number of particles thermal systems consist of, and on the other the more traditional ergodic approach. These two contributions are expressed in two distinct ingredients in the scheme which Malament and Zabell offer to explain the success of the microcanonical phase averages. Both ingredients are incapable of explaining the success of the averaging method on their own, but the combination constitutes an explanation scheme that is watertight.

The first claim is the validity of what Malament and Zabell call the Khinchin–Lanford dispersion theorems (Khinchin, 1949; Lanford, 1973). In their general form, these limit theorems show that with large microcanonical probability, phase functions are always close to their microcanonical averages. What is still lacking, however, is a translation of the phrase 'large microcanonical probability'.

The second claim is the statement that the 'microcanonical measure actually represents the probability of finding an isolated equilibrium system in a particular microstate' (Malament and Zabell, 1980, p. 343). Taken by itself, this claim is not sufficient for the explanatory purposes either. Although it says nothing less than that the microcanonical measure is the appropriate probability measure, it does not explain why observed values are always equal to phase averages and are not spread around those averages. However, combining the two claims does result in an explanation of the success of the microcanonical phase averaging procedure. This is because the combination leads to the statement that the *actual* probability that phase functions are always close to their microcanonical averages is large.

How do Malament and Zabell justify the second claim? This is where ergodic theory comes in. They use another corollary of the ergodic theorem, which says that for metrically transitive systems the microcanonical measure is the unique measure which is both stationary and absolutely continuous. They argue that the 'actual' probability measure (which I will denote $P_{@}$) for an isolated system in equilibrium should be both stationary and absolutely continuous. Thus, for metrically transitive systems at least, they come to the conclusion that the 'actual' probability measure is equal to the microcanonical measure.

They also prove a theorem which demonstrates the equivalence of absolute continuity and a property called 'translation continuity' (continuity of the probability distribution under translation of sets in phase space). The importance of this theorem is that one may now try to justify the translation continuity of $P_{@}$ rather than its absolute continuity; this might be done by appealing either to stability under actual perturbations or to measurement imprecision (see Van Lith (2001) for an extensive discussion).

The distinction between the 'actual probability' $P_{@}$ and the microcanonical probability measure μ is conceptually important. Malament and Zabell attribute certain properties to $P_{@}$, namely stationarity and translation continuity. They claim that these properties are natural and plausible. The interpretation of probability comes into play when evaluating whether these assumptions are indeed plausible. In contrast to $P_{@}$, the microcanonical measure μ should be viewed as a mathematical entity which is not *a priori* associated with the actual probability. The properties of μ can be studied with the tools provided by mathematical physics; especially, the question whether the dynamical system $\langle \Gamma, \mathscr{B}, \mu, T_t \rangle$ is ergodic can be answered in that way. The separation between the notion of actual probability or chance on the one hand and the descriptive apparatus in the form of a probability measure on the other is in my view an illuminating and important contribution to the debate.

But there are several weaknesses of this approach. First of all, the approach is limited to metrically transitive systems, so that the criticism by Earman and Rédei again applies. Indeed, for systems which have only small regions of non-ergodicity (e.g. systems with small KAM-tori) the approach does not have anything to say. But here a modification of Malament and Zabell's work by Vranas offers a way out; his approach also applies to systems that are 'ε-ergodic', which, roughly, means that small regions of non-ergodicity are allowed (Vranas, 1998). Secondly, more has to be said about a justification of both conditions that are imposed on $P_{@}$ (stationarity, and absolute or translation continuity).

To conclude, I think this approach is a valuable, but not yet fully developed, alternative to the standard ergodic approach: it certainly deserves further study.

4. Stationarity

4.1. Does stationarity represent equilibrium?

As an illustration of the differences between the ergodic approaches outlined above, and of the importance of the interpretations of probability, let us look at the notion of stationarity in more detail. It is clear that stationarity plays an important role in all three ergodic approaches. In Malament and Zabell's scheme stationarity is one of the main assumptions. The ergodic decomposition theorem is a representation theorem for stationary measures. Finally, time averages are stationary by definition.

In all cases stationarity of the probability distribution is somehow connected to the system being in equilibrium. Indeed, Malament and Zabell reflect the usual standpoint when they write:

> whatever one's account of probability, it is fundamental on the Gibbsian view that equilibrium probabilities are stationary (Malament and Zabell, 1980, p. 345).

But, I submit, there is a problem with the standard Gibbsian account of equilibrium as a stationary probability distribution, which suggests that this account should be weakened. Also, as I will argue below, it depends on one's interpretation of probability how stringent this problem is. In the following I will present an alternative account of equilibrium, and investigate the consequences for the three ergodic approaches discussed above.

The problem is that equating equilibrium with stationarity of the probability distribution makes it impossible that for an isolated system, a non-equilibrium state will evolve into an equilibrium state by the equations of motion of classical mechanics. This is in striking contrast with thermodynamics, because thermodynamics clearly demands such evolutions. Thus, a statistical mechanical explanation in mechanical terms of such thermal processes as the approach to equilibrium is blocked.

A similar point has been raised in Leeds (1989) as a criticism of Malament and Zabell's paper. Malament and Zabell conclude from the assumptions they make about the probability distribution that the probability is uniquely given by the microcanonical distribution. But then, Leeds argues, it can never have been different from the microcanonical distribution, since that distribution is stationary under the equations of motion. But surely the situation prior to equilibrium may have been such that the microcanonical distribution does not properly give the probabilities on phase space. Therefore, the distribution that represents equilibrium cannot be strictly stationary.

Thus, the point is that, with equilibrium represented as a stationary probability distribution, no transition from non-equilibrium to equilibrium is possible governed by the Hamiltonian equations of motion. But now the interpretation of probability becomes relevant. On a subjective interpretation of probability (either personalist, or 'inter-subjective' as for instance in Jaynes's Maximum Entropy formalism) probabilities represent degrees of belief, or they characterise the available information about the system. Therefore, probabilities evolve not only under the influence of the equations of motion, but also when new information comes in, or when a person changes his beliefs. This means that the fact that transitions from non-stationary to stationary distributions are not allowed by the dynamics is especially a problem for objective interpretations of probability.

As a solution to this problem (a solution which, for the reason just given, applies especially to the frequency interpretation of probability), I propose the following definition (see also Van Lith, 1999, 2001):

Definition 1. (**ε-Equilibrium**). *Let a class of functions* $\Omega \ni F$ *be given, and let* $\varepsilon = \{\varepsilon_F\}$ *be fixed. Then a system is in ε-equilibrium during the time interval* τ *iff* $\forall F \in \Omega \; \exists c_F \; \forall t \in \tau \; |\langle F \rangle_{P_t} - c_F| \leqslant \varepsilon_F$.

In words: during the time interval τ the expectation values of functions in the class Ω are time-independent, or may fluctuate in time at most within some small, fixed intervals ε_F. Irrespective of the exact choice of the class Ω and of the ε_F, dynamical evolution from a non-equilibrium state to an equilibrium state is now possible, as long as the ε_F are non-zero.

4.2. The role of stationarity in the ergodic approaches

Let us now turn to the question of how the observation that non-stationary distributions cannot evolve into stationary distributions affects the three ergodic approaches. In other words, the question is whether the three approaches can be reconciled with the thermodynamic fact that transitions between non-equilibrium and equilibrium states do occur in nature.

The time average interpretation of probabilities does not possess the means to discriminate between equilibrium and non-equilibrium. Probabilities are stationary by definition. This follows straightforwardly from their definition as expressed in relation (1). Thus, no embedding of equilibrium theory in the general framework of non-equilibrium statistical mechanics is possible in which this interpretation is pursued consistently. Transitions from a non-equilibrium state to an equilibrium state cannot be accounted for in this manner. This renders the interpretation untenable, in my view.

The ergodic decomposition theorem, which supports the personalistic interpretation of probability, is a representation theorem valid for stationary probability measures. Guttmann discusses four possible options the subjectivist has with which to justify stationarity. However, these options, as Guttmann shows, all fail (Guttmann, 1999). This is because to a subjectivist only coherence requirements are compelling, but stationarity cannot be enforced on grounds of coherence. But the same can be said of exchangeability in De Finetti's original theorem and yet the representation theorem is seen as very important in that context.

Beliefs can change from person to person, and there are no coherence requirements that fix a particular characterisation of equilibrium, such as stationarity. This is yet another reason why proponents of the personalist interpretation need not be bothered at all by the fact that the Hamiltonian equations of motion do not allow for transitions from non-stationary to stationary probability distributions.

In Malament and Zabell's ergodic approach the fact that the measure is *strictly* stationary in time is crucial. This is because they offer a uniqueness theorem; if measures that are *approximately* stationary were allowed, the microcanonical measure would lose its special status, and the explanatory scheme would lose its force. Thus, the observation that equilibrium should be represented by a probability distribution that is only approximately stationary is a severe problem for this approach.

Here again the modification made by Vranas offers help. Although his theorems are derived using strict stationarity, he writes in a footnote that this can be weakened to a notion he coins 'forward ε-stationarity' (Vranas, 1998, p. 702, footnote 12). This, roughly, means that for each measurable set A the probability $P(A)$ is now allowed to fluctuate within small bounds, analogously to ε-equilibrium as defined above. Strict stationarity would demand constant probabilities. Also with this notion transitions are possible from a state which is not forward ε-stationary to a state which is.

References

Earman, J. and Rédei, M. (1996) 'Why Ergodic Theory Does Not Explain the Success of Equilibrium Statistical Mechanics', *British Journal for Philosophy of Science* **47,** 63–78.

Ehrenfest, P. and Ehrenfest-Afanassjewa, T. (1912) *The Conceptual Foundations of the Statistical Approach in Mechanics* (Ithaca, NY: Cornell University Press, 1959; New York: Dover, 1990).

Guttmann, Y. M. (1999) *The Concept of Probability in Statistical Physics* (Cambridge: Cambridge University Press).

Jaynes, E. T. (1967) 'Foundations of Probability Theory and Statistical Mechanics', in M. Bunge (ed.), *The Delaware Seminar in the Foundations of Physics*, Vol. 1 (New York: Springer-Verlag).

Khinchin, A. I. (1949) *Mathematical Foundations of Statistical Mechanics* (New York: Dover).

Lanford, O. E. (1973) 'Entropy and Equilibrium States in Classical Statistical Mechanics', in A. Lenard (ed.), *Statistical Mechanics and Mathematical Problems*, Springer Lecture Notes in Physics (Berlin: Springer-Verlag).

Leeds, S. (1989) 'Discussion: Malament and Zabell on Gibbs Phase Averaging', *Philosophy of Science* **56,** 325–340.

Malament, D. B. and Zabell, S. L. (1980) 'Why Gibbs Phase Averages work—The Role of Ergodic Theory', *Philosophy of Science* **47,** 339–349.

Sklar, L. (1973) 'Statistical Explanation and Ergodic Theory', *Philosophy of Science* **40,** 194–212.

Sklar, L. (1993) *Physics and Chance. Philosophical Issues in the Foundations of Statistical Mechanics* (Cambridge: Cambridge University Press).

Van Lith, J. (1999) 'Reconsidering the Concept of Equilibrium in Classical Statistical Mechanics', *Philosophy of Science* **66** (Proceedings), S107–S118.

Van Lith, J. (2001) 'Stir in Stillness. A Study in the Foundations of Equilibrium Statistical Mechanics', Ph.D. thesis, Utrecht University.

Von Mises, R. (1931) *Wahrscheinlichkeitsrechnung und ihre Anwendung in der Statistik und theoretischen Physik* (New York: Rosenberg).

Von Plato, J. (1988) 'Ergodic Theory and the Foundations of Probability', in B. Skyrms and W. L. Harper (eds), *Causation, Chance, and Credence*, Proceedings of the Irvine Conference on Probability and Causation, Vol. 1 (Dordrecht: Kluwer Academic), pp. 257–277.

Von Plato, J. (1989) 'Probability in Dynamical Systems', in J. E. Fenstad, I. T. Frolov and R. Hilpinen (eds), *Logic, Methodology and Philosophy of Science VIII*, Proceedings of the Eighth International Congress of Logic, Methodology and Philosophy of Science, Moscow, 1987 (Amsterdam: North Holland).

Vranas, P. B. M. (1998) 'Epsilon-Ergodicity and the Success of Equilibrium Statistical Mechanics', *Philosophy of Science* **65,** 688–708.

Wightman, A. S. (1985) 'Regular and Chaotic Motions in Dynamical Systems: Introduction to the Problems', in G. Velo and A. S. Wightman (eds), *Regular and Chaotic Motions in Dynamical Systems* (New York: Plenum Press), pp. 1–26.