Chapter 2
Probability and Time Symmetry in Classical Markov Processes

Guido Bacciagaluppi

2.1 Introduction

The problem of the arrow of time in physics is that certain phenomena appear systematically to take place much more frequently than their time reversals, and this despite the fact that the fundamental laws are mostly believed to be fully time-symmetric, at least as long as they are deterministic. The two common general strategies for addressing this problem use, respectively, time-asymmetric laws or time-symmetric laws with special initial or boundary conditions.

It is less clear that such a problem exists also if one assumes indeterministic laws, since, intuitively, probabilities may be thought of as intrinsically time-directed. However, one should distinguish sharply between issues in the interpretation of probability, where these intuitions are strongest ('open future' versus 'fixed past'), and issues of formalism, which are the only ones involved in the description of the phenomena (can time-directed behaviour be described by formally time-symmetric laws?).

In this paper we propose to investigate, in the simple abstract setting of discrete Markov processes (more precisely, Markov processes with discrete state space and continuous time), whether and in what sense time-directed behaviour might indeed be compatible with time-symmetric probabilistic laws. We shall argue that time-symmetric stochastic processes, in a classical setting, are indeed quite capable of describing time-directed behaviour (or, when otherwise, that the remaining time asymmetry is quite benign). Thus, we suggest that a move to indeterministic laws is not likely to change the terms of the debate on the arrow of time. There will still be two fundamental alternatives for describing time-directed behaviour: adopting time-asymmetric laws, or adopting time-symmetric laws and suitable boundary conditions.[1]

G. Bacciagaluppi (✉)
Department of Philosophy, University of Sydney, Sydney, NSW 2006, Australia
e-mail: guido.bacciagaluppi@arts.usyd.edu.au

[1]Note that Markov processes are indeed sometimes used in the context of thermodynamics to explain the thermodynamic arrow in terms of a 'probabilistic arrow of time'. Uffink (2007,

M. Suárez (ed.), *Probabilities, Causes and Propensities in Physics*, Synthese Library 347, DOI 10.1007/978-1-4020-9904-5_2, © Springer Science+Business Media B.V. 2011

On the basis of these results we then argue that considering the arrow of time in a probabilistic setting fails to justify a qualitative distinction in status between the future and the past. Of course, investigating notions of time symmetry or asymmetry at the level of the formalism can yield no normative conclusion about the interpretation of probability. However, we take it that it can provide useful guidelines for choosing or constructing a good interpretation, and in this sense we suggest that the common interpretation of probabilities as time-directed is unjustified.

Our results apply to classical probabilities. In a separate paper (Bacciagaluppi, 2007), we discuss the case of quantum probabilities as they appear in no-collapse approaches to quantum mechanics, specifically in the context of the decoherent histories formalism of quantum mechanics. The conclusions drawn in the two papers are quite different. Whereas in the classical case we shall argue against drawing such distinctions, in the quantum case we find that, albeit in a restricted sense, a qualitative distinction between forwards and backwards probabilities can be justified.

The structure of this paper is as follows: after reviewing some elementary theory in Section 2.2, we shall discuss notions of time symmetry for discrete Markov processes in Section 2.3. Then, in Section 2.4, we shall review reasons given for a time-asymmetric treatment of probabilities (Section 2.4.1); argue that, contrary to appearances, the relevant examples can very well be treated using processes that are time-symmetric or only harmlessly time-asymmetric (Section 2.4.2); and, finally, draw lessons for the interpretation of probability (Section 2.4.3).

2.2 A Few Essentials About Markov Processes

A stochastic process is defined to be a family of random variables, indexed by t, from a probability space Ω to a (common) state space S, which for the purposes of this paper we shall assume to be discrete (and sometimes finite):

$$X(t,.): \quad \Omega \to S\,. \tag{2.1}$$

It is, however, simpler to discuss a stochastic process in terms of joint distributions at finitely many times. Indeed, a classic theorem by Kolmogorov (1931) states that a stochastic process can be reconstructed from the collection of its *finite-dimensional distributions*, the n-fold joint distributions for all n:

$$p_{i_1 i_2 \dots i_n}(t_1, t_2, \dots, t_n)\,. \tag{2.2}$$

We shall also assume that the process is Markov, i.e. for any $t_1 < t_2 < \dots < t_j < t_{j+1} < \dots < t_n$,

Section 7) has independently criticised such attempts in a way that is very close to the ideas expressed in this paper.

$$p_{i_{j+1}\dots i_n|i_1\dots i_j}(t_{j+1},\dots,t_n|t_1,\dots,t_j) = p_{i_{j+1}\dots i_n|i_j}(t_{j+1},\dots,t_n|t_j)\ , \tag{2.3}$$

i.e.

$$\frac{p_{i_1\dots i_n}(t_1,\dots,t_n)}{p_{i_1\dots i_j}(t_1,\dots,t_j)} = \frac{p_{i_j\dots i_n}(t_1,\dots,t_n)}{p_{i_j}(t_n)}\ . \tag{2.4}$$

The finite-dimensional distributions of a Markov process can be reconstructed from its two-dimensional distributions,

$$p_{ij}(t,s)\ , \tag{2.5}$$

as is easily shown by induction. It should also be noted that the Markov condition is only apparently time-directed. Indeed, (2.4) is equivalent to

$$\frac{p_{i_1\dots i_n}(t_1,\dots,t_n)}{p_{i_j\dots i_n}(t_j,\dots,t_n)} = \frac{p_{i_1\dots i_j}(t_1,\dots,t_j)}{p_{i_j}(t_j)}\ , \tag{2.6}$$

i.e.

$$p_{i_1\dots i_{j-1}|i_j\dots i_n}(t_1,\dots,t_{j-1}|t_j,\dots,t_n) = p_{i_1\dots i_{j-1}|i_j}(t_1,\dots,t_{j-1}|t_j)\ , \tag{2.7}$$

so that the Markov condition is itself still perfectly time-symmetric.

Now we can introduce (two-time) transition probabilities. That is, for $t > s$ we define:

$$p_{i|j}(t|s) := \frac{p_{ij}(t,s)}{p_j(s)} \tag{2.8}$$

(*forwards* transition probabilities), and

$$p_{i|j}(s|t) := \frac{p_{ij}(s,t)}{p_j(t)} = \frac{p_{ji}(t,s)}{p_j(t)} \tag{2.9}$$

(*backwards* transition probabilities).

Using the forwards transition probabilities we can express the time evolution of the single-time distributions as

$$p_i(t) = \sum_j p_{i|j}(t|s)p_j(s)\ , \tag{2.10}$$

which we can also write in more compact form as

$$\mathbf{p}(t) = P(t|s)\mathbf{p}(s)\ . \tag{2.11}$$

$P(t|s)$ is called the transition matrix, mapping the probability vector $\mathbf{p}(s)$ into $\mathbf{p}(t)$. The matrix $P(t|s)$ is a so-called stochastic matrix, i.e. all elements of $P(t|s)$ are between 0 and 1, and each column of $P(t|s)$ sums to 1.

Similarly, we have the time-reversed analogues of (2.10) and (2.11):

$$p_i(s) = \sum_j p_{i|j}(s|t)p_j(t) , \tag{2.12}$$

and

$$\mathbf{p}(s) = P(s|t)\mathbf{p}(t) . \tag{2.13}$$

Note that the backwards transition matrix $P(s|t)$ is not in general the inverse matrix $P(t|s)^{-1}$, as can be seen easily by noting that the former is always well-defined, via (2.9), but the latter is not: e.g. if for given t and s,

$$P(t|s) = \begin{pmatrix} 1-\varepsilon & \alpha \\ \varepsilon & 1-\alpha \end{pmatrix} , \tag{2.14}$$

AQ1 invertibility rules out the case $\alpha = 1 - \varepsilon$.

The intuitive reason for this discrepancy is that, given (2.8) and (2.9), $\mathbf{p}(s)$ is not in general specifiable independently of both $P(t|s)$ and $P(s|t)$. Therefore, the condition that for all s and t,

$$\mathbf{p}(s) = P(s|t)P(t|s)\mathbf{p}(s) , \tag{2.15}$$

does not imply

$$P(s|t)P(t|s) = \mathbf{1} , \tag{2.16}$$

because $\mathbf{p}(s)$ in (2.15) is not arbitrary.

Now, let us take two possibly different initial distributions and evolve them both in time using the same (forwards) transition probabilities. It is then elementary to show that

$$\begin{aligned} \sum_i |p_i(t) - q_i(t)| &= \sum_i \Big| \sum_j p_{i|j}(t|s)p_j(s) - \sum_j p_{i|j}(t|s)q_j(s) \Big| \\ &\leq \sum_i \sum_j |p_{i|j}(t|s)||p_j(s) - q_j(s)| \\ &= \sum_j |p_j(s) - q_j(s)| . \end{aligned} \tag{2.17}$$

It follows that $\sum_i |p_i(t) - q_i(t)|$ converges to some positive number, not necessarily zero. Under suitable conditions, in particular if there are 'enough' transitions, one can hope to strengthen this result to

$$\lim_{t \to \infty} \sum_j |p_j(t) - q_j(t)| = 0 \,, \tag{2.18}$$

i.e. any two distributions would converge asymptotically. Under appropriate conditions, there would even be convergence of any initial distribution towards a unique (time-independent) limit distribution.

'Limit theorems', or 'ergodic theorems' for discrete Markov processes describe precisely the asymptotic properties of processes with a given set of (forwards) transition probabilities, in particular the circumstances under which such processes converge to a limit (uniquely or non-uniquely), and the relevant notion and corresponding speed of convergence. Analogous results hold, of course, if one fixes the set of backwards transition probabilities.[2]

Let us define state j to be a *consequent* of state i, if for all times s with $p_i(s) \neq 0$ there is a $t > s$ such that $p_{j|i}(t|s) \neq 0$. A state i is *transient* iff there is a state j that is a consequent of i, but such that i is not a consequent of j. The relation of consequence defines equivalence classes on the non-transient states (so-called ergodic classes).

In the case of finitely many states a sufficient condition for the existence of an (invariant) *limit distribution* for $t \to \infty$ is that the (forwards) transition probabilities are time-translation invariant – synonyms: if the (forwards) transition probabilities are stationary, or if the process is (forwards) *homogeneous*. The limit distribution decomposes into a convex combination of the limit distributions on each ergodic class, while the probability of any transient state converges to zero (see e.g. Doob, 1953, Chapter VI). In the next section and the appendix, we shall need to refer to the case of discrete time, where the result is slightly weaker, since in some ergodic classes one may have cyclic behaviour rather than convergence (see e.g. Doob, 1953, Chapter V).

Returning to the case of continuous time, if one has denumerably many states, homogeneity is not sufficient for the existence of limit distributions, and additional conditions can be used. On the other hand, homogeneity is not a necessary condition either for the existence of limit distributions, and alternative sufficient conditions are known. As an example, take a two-state process that has equal probabilities for jumping from 0 to 1 as from 1 to 0 in any given time interval, and such that in unit time these probabilities are always larger than a given δ. Then one can easily see that the process will converge exponentially fast towards the invariant distribution $p_0(t) = p_1(t) = 1/2$, whether or not the transition probabilities are time translation invariant. Similarly, there are conditions that ensure asymptotic convergence when the process has no invariant limit distribution (see e.g. Hajnal, 1958).

If the single-time distribution $p_i(t)$ of a process is invariant, it is itself equal to the limit distribution of the process, and we shall say that the process is in *equilibrium*.

[2]For a good introduction to the complex theme of ergodic theory in the deterministic case, see Uffink (2007, Section 6).

(We shall occasionally also refer to an invariant distribution as an equilibrium distribution.) Note that if a process is in equilibrium, it has no transient states. Finally, a process that is both homogeneous and in equilibrium is said to be *stationary*.

2.3 Definitions of Time Symmetry

The framework we have introduced above is quite austere, and we must realise that, at least for the purpose of investigating time symmetry, it has its limitations. For instance, we do not have enough structure to define the time reverse of a state (there is no analogue of inversion of momenta in Newtonian mechanics, for instance). More importantly, we are not going to be able to identify and abstract from *systematic* components of the process, in particular components that may appear time-asymmetric but might in fact be generated by some time-symmetric law (think of a diffusion process taking place in a Newtonian gravitational field). Nevertheless, the insights we shall gain will be enough to discuss how typical examples of time-directed behaviour can be described in terms of time-symmetric processes, and to provide clues as to the time-symmetric or time-asymmetric status of the probabilities with respect to their interpretation.

It is natural to consider transition probabilities as what defines the dynamics of a system described by a Markov process. This in turn suggests to consider the following condition as a possible condition for a time-symmetric process: that forwards and backwards transition probabilities coincide, i.e. (for all i, j, t and s)

$$p_{i|j}(t|s) = p_{i|j}(s|t) \tag{2.19}$$

or (for all t, s)

$$P(t|s) = P(s|t) \; . \tag{2.20}$$

This is by analogy to the condition, familiar from the deterministic case, that the backwards equations of motion have the same form as the forwards equations.

In the literature on Markov processes, however, the usual condition of time symmetry is the so-called condition of *detailed balance*[3]:

$$p_{i|j}(t|s)p_j(s) = p_{j|i}(t|s)p_i(s) \; . \tag{2.21}$$

The meaning of detailed balance can be readily seen using the notion of *probability current*, i.e. the net probability flow from a state j to a state i between s and t:

$$j_{ij}(t,s) := p_{i|j}(t|s)p_j(s) - p_{j|i}(t|s)p_i(s) \; . \tag{2.22}$$

[3] My thanks to Werner Ehm for discussions about this notion.

Detailed balance simply means that there are no probability currents.

Our main purpose in this section will be to see that the two conditions (2.19) and (2.21) are equivalent, at least under certain conditions. Note that (2.21) is often formulated under the additional presupposition that the process is stationary, but we shall *not* make this assumption.

Symmetry of the transition probabilities obviously involves both forwards and backwards transition probabilities, while detailed balance explicitly involves only the forwards transition probabilities. On the other hand,

$$\begin{aligned} j_{ij}(t,s) &= p_{ij}(t,s) - p_{ji}(t,s) \\ &= p_{ij}(t,s) - p_{ij}(s,t) , \end{aligned} \tag{2.23}$$

therefore detailed balance is equivalent to symmetry of the two-time distributions,

$$p_{ij}(t,s) = p_{ij}(s,t) , \tag{2.24}$$

which is clearly a time symmetry condition.

Now, (2.24) and hence detailed balance are easily seen to be a sufficient condition for both equilibrium and the symmetry of transition probabilities (2.19). Indeed, performing a sum over i in (2.24) yields invariance of the single-time distributions:

$$p_j(s) = p_j(t) , \tag{2.25}$$

i.e. equilibrium. But from (2.24) and (2.25) we obtain:

$$p_{i|j}(t|s) = \frac{p_{ij}(t,s)}{p_j(s)} = \frac{p_{ij}(s,t)}{p_j(s)} = \frac{p_{ij}(s,t)}{p_j(t)} = p_{i|j}(s|t) , \tag{2.26}$$

i.e. (2.19), as long as either side is well-defined.

Notice that, conversely, (2.19) and equilibrium together imply (2.24) and therefore detailed balance. Indeed,

$$\begin{aligned} p_{ij}(t,s) &= p_{i|j}(t|s)p_j(s) \\ &= p_{i|j}(s|t)p_j(s) \\ &= p_{i|j}(s|t)p_j(t) = p_{ij}(s,t) . \end{aligned} \tag{2.27}$$

Instead, equilibrium on its own does not imply detailed balance (and therefore not symmetry of transition probabilities either). Indeed, take a three-state system with

$$P(t|s) = \begin{pmatrix} 1/3 & 1/6 & 1/2 \\ 1/2 & 1/3 & 1/6 \\ 1/6 & 1/2 & 1/3 \end{pmatrix}^{t-s} . \tag{2.28}$$

We have in particular that

$$
\begin{aligned}
p_{i|i}(t+1|t) &= 1/3 \, , \\
p_{i+1|i}(t+1|t) &= 1/2 \, , \\
p_{i-1|i}(t+1|t) &= 1/6
\end{aligned}
\tag{2.29}
$$

(where $i+1$ and $i-1$ are to be read as addition mod 3). The equilibrium distribution for this process is $p_i(t) = 1/3$, but there is clearly a non-zero current $0 \to 1 \to 2 \to 0$, and detailed balance fails.

This example is generic in the sense that the only way to have currents in equilibrium, whether for finite or denumerable state space, is to have a circular current, i.e. a current along a closed chain of states with at least three members,[4]

$$
i \to j \to k \to i \, . \tag{2.30}
$$

Therefore equilibrium and zero circular currents together are equivalent to detailed balance. In the special case of a two-state system, there are no three-element chains, and equilibrium is in fact equivalent to detailed balance.

Simple examples suggest that, under suitable conditions, symmetry of the transition probabilities (2.19) might in fact imply equilibrium and therefore (by (2.27)) be equivalent to detailed balance. Take a homogeneous two-state process with (forwards) transition matrix

$$
P(t|s) = \begin{pmatrix} 1-\alpha & \varepsilon \\ \alpha & 1-\varepsilon \end{pmatrix}^{t-s} \, . \tag{2.31}
$$

If we take $\alpha \neq 0$ and ε arbitrary, this is a toy model of decay (with non-zero probability α of decay in unit time), with or without re-excitation (depending on whether $\varepsilon \neq 0$ of $\varepsilon = 0$).

Imposing (2.19) in this example leads to

$$
p_0(t) = \frac{\alpha}{\alpha+\varepsilon} \, , \qquad p_1(t) = \frac{\varepsilon}{\alpha+\varepsilon} \tag{2.32}
$$

for all t, i.e. the single-time distribution is fully constrained to be the equilibrium distribution of the process (and the process is stationary).

[4] In the case of denumerable state space, assume there are non-zero currents in equilibrium but no circular currents. Let us say that, between s and t, state 0 gains probability ε from states $1, \ldots, i_1$ (distinct from 0). Obviously, $\sum_{i=1}^{i_1} p_i(s) \geq \varepsilon$. In the same time interval, the states $1, \ldots, i_1$ must gain probability at least ε from some states $i_1+1, \ldots, i_2$ (all distinct from $0, \ldots, i_1$), and $\sum_{i=i_1+1}^{i_2} p_i(s) \geq \varepsilon$. Therefore $\sum_{i=1}^{i_2} p_i(s) \geq 2\varepsilon$. Repeat the argument until $\sum_{i=1}^{i_n} p_i(s) \geq n\varepsilon > 1$, which is impossible.

Indeed, for arbitrary t and s define α_{t-s} and ε_{t-s} such that

$$P(t|s) = \begin{pmatrix} 1-\alpha_{t-s} & \varepsilon_{t-s} \\ \alpha_{t-s} & 1-\varepsilon_{t-s} \end{pmatrix} . \tag{2.33}$$

Then, from

$$p_{0|1}(t|s) = p_{0|1}(s|t) = \alpha_{t-s} \tag{2.34}$$

and

$$p_{1|0}(t|s) = p_{1|0}(s|t) = \varepsilon_{t-s} , \tag{2.35}$$

one obtains

$$\varepsilon_{t-s} p_0(s) = \alpha_{t-s} p_1(t) , \tag{2.36}$$

$$\varepsilon_{t-s} p_0(t) = \alpha_{t-s} p_1(s) . \tag{2.37}$$

Thus, since there are only two states,

$$\varepsilon_{t-s} p_0(s) = \alpha_{t-s}(1 - p_0(t)) , \tag{2.38}$$

$$\varepsilon_{t-s} p_0(t) = \alpha_{t-s}(1 - p_0(s)) , \tag{2.39}$$

whence

$$p_0(s) = p_0(t) = \frac{\alpha_{t-s}}{\alpha_{t-s} + \varepsilon_{t-s}} . \tag{2.40}$$

Therefore, $p_0(t)$ is constant, since t and s are arbitrary. Finally, substituting $s = t-1$ in (2.40), we have

$$p_0(t) = \frac{\alpha}{\alpha + \varepsilon} , \tag{2.41}$$

and the claim follows.

We now ask for conditions under which symmetry of the transition probabilities strictly implies equilibrium and thus becomes equivalent to detailed balance.

Let us first specialise to homogeneous Markov processes, i.e. the transition probabilities are time-translation invariant. Then equilibrium follows very easily. (Incidentally, note that a forwards or backwards homogeneous process satisfying (2.19) will be both forwards and backward homogeneous.) Indeed, for all t, s,

$$\mathbf{p}(t+s) = P(t+s|t+s/2)P(t+s/2|t)\mathbf{p}(t) . \tag{2.42}$$

By translation invariance,

$$\mathbf{p}(t+s) = P(t+s/2|t)P(t+s/2|t)\mathbf{p}(t) , \tag{2.43}$$

and by symmetry

$$\mathbf{p}(t+s) = P(t|t+s/2)P(t+s/2|t)\mathbf{p}(t)\ , \tag{2.44}$$

but by definition also

$$\mathbf{p}(t) = P(t|t+s/2)P(t+s/2|t)\mathbf{p}(t)\ . \tag{2.45}$$

Therefore,

$$\mathbf{p}(t+s) = \mathbf{p}(t) \tag{2.46}$$

for all t, s, i.e. the process is in equilibrium.

If we relax the assumption that the process is homogeneous, it is still a theorem that (2.19) implies equilibrium, at least under the further assumptions that (a) the state space has finite size n, and (b) for all i, j and s the transition probabilities $p_{i|j}(t|s)$ are continuous in t. (The appendix provides an elementary derivation of this result from the ergodic theorem for discrete time.) Thus, under the appropriate conditions, the two definitions of time symmetry (2.19) and (2.21) are indeed equivalent.

2.4 Probability and Time Symmetry

2.4.1 Arguments for Asymmetry

Imagine a world in which fundamental laws are probabilistic. Imagine further that this world contains an arrow of time, that is, typical examples of time-directed behaviour, and that this behaviour is investigated by observers who can set up experiments under controlled initial conditions (but not final ones). That is, like ourselves, observers in this world are subject to some macroscopic arrow of time that may or may not be related to the time-directed behaviour under scrutiny. Finally, let this be a classical world; in particular, assume that gaining knowledge of the state i of a system at a certain time (in particular with regard to alternative initial conditions) can be done in principle without disturbing the system, so that we can still consider it as governed by the same stochastic process.

It will be tempting to interpret the probabilistic laws in this world as intrinsically time-directed. Such laws will specify objective probabilities for events in the future given events in the present (if the laws are Markovian), while probabilities for past events will be regarded as merely epistemic. The underlying intuition is that, under indeterminism, the future is genuinely 'open', while the past, while perhaps unknown, is 'fixed'.

Formally, however, there is a very good argument for saying that in a classical stochastic process there is no distinction between future and past: a classical stochastic process is defined as a probability measure over a space of trajectories, so the formal definition is completely time-symmetric. Transition probabilities towards the future can be obtained by conditionalising on the past; but, equally, transition probabilities towards the past can be obtained by conditionalising on the future. Individual trajectories may exhibit time asymmetry, and there may be a quantitative asymmetry between forwards and backwards transition probabilities, but at least as long as the latter are not all 0 or 1, quantitative differences fall short of justifying a notion of fixed past.

On the other hand, at least in a world as the one sketched above, there are ways of arguing for *qualitative* formal differences between forwards and backwards transitions probabilities that could suggest also a different interpretational status for the two kinds of probabilities:

(A) In a probabilistic setting one has good ergodic behaviour, in particular, if time translation invariance of the transition probabilities holds (assuming finiteness of the state space or other suitable conditions), one will have a tendency for a stochastic process to equilibrate in time, regardless of the initial distribution. Such an arrow of time would thus appear to be very deeply seated in the use of probabilistic concepts. A related argument is that in the homogeneous case (and, as we have mentioned, more generally) the symmetry of transition probabilities implies equilibrium, and thus rules out not only any equilibration process but any time development of the probabilities whatsoever (Sober, 1993).

(B) Another interesting argument for asymmetry between forwards and backwards probabilities runs along the following lines. Take the simple model of exponential decay (2.31), with probability α of decay from the excited state 1 to the ground state 0 in unit time, and starting with all 'atoms' excited, i.e. . We have:

$$p_{0|1}(t+1|t) = \alpha \ , \tag{2.47}$$

for all t, but:

$$p_{0|1}(t|t+1) = \begin{cases} \to \alpha \ \text{for}\ t \to \infty \ , \\ \to 0 \ \text{for}\ t \to 0 \ . \end{cases} \tag{2.48}$$

In this example, the forwards transition probabilities are time translation invariant, but the backwards transition probabilities are not. This difference has been used to argue that forwards transition probabilities are indeed law-like, while backwards transition probabilities are epistemic (Arntzenius, 1995).

(C) Finally, backwards transition probabilities are not invariant across experiments when one varies the initial distribution. One can thus argue that if the initial distribution of the process is an epistemic distribution over contingent initial states, then the backwards transition probabilities cannot be law-like, or not entirely law-like, because they depend on the epistemic initial distribution. A related argument is that, in general, at most one set of transition probabilities can be law-like, otherwise also the single-time probabilities will be, so that it appears that initial conditions cannot be freely chosen (Watanabe, 1965, Section 5).

These arguments infer from typical time-directed behaviour to formal qualitative differences in the transition probabilities. It is this type of inference that we shall question below. Without a qualitative difference in the formalism, however, we take it that there is no reason to deny the same interpretational status to both sets of transition probabilities alike.

2.4.2 *Time-Directed Behaviour and Time-Symmetric Probabilities*

The situation of convergence to equilibrium – indeed, the simple example of decay – can be used to exemplify at once all three purported differences between forwards and backwards transition probabilities and, at least at first sight, seems thus to be totally intractable in terms of symmetric processes. Indeed, (A) we have seen that time symmetry of transition probabilities implies equilibrium of the process ((2.32) above). (B) We have also seen the lack of time translation invariance for the backwards transition probabilities ((2.47) and (2.48) above). Finally, (C) if we start with all 'atoms' in the ground state, i.e. $p_0(0) = 1$, we obtain:

$$p_{0|1}(t+1|t) = \alpha \ , \tag{2.49}$$

for all t, but:

$$p_{0|1}(t|t+1) = \begin{cases} \to \alpha \ \text{for} \ t \to \infty \ , \\ \to 1 \ \text{for} \ t \to 0 \ . \end{cases} \tag{2.50}$$

Thus, a different choice of initial condition will indeed lead to different backwards transition probabilities.

The question we wish to raise is: can we indeed infer that there are such differences in the transition probabilities from time asymmetries of the phenomena, i.e. from the time-directed behaviour of *samples*?

Obviously, one must distinguish between the transition *probabilities* of the process and the transition *frequencies* in any actual sample. Observed behaviour, in particular time-directed behaviour, will always be defined in terms of frequencies, and in order to conclude from frequencies to probabilities, we have to ensure that the sample is unbiased. Indeed, suppose that we bias the sample by performing a post-selection of the final ensemble. Then in general we shall influence the *forwards* transition frequencies, in particular destroying their time translation invariance.

If we assume that the process has a limit distribution for $t \to \infty$, a simple criterion to make sure that the final ensemble is sufficiently unbiased is to check whether the distribution of the sample is at least approximately time-independent, i.e. whether or not the sample has been prevented from equilibrating or has subsequently departed from equilibrium for any reason (a statistical fluctuation, a final cause, or an uncooperative lab assistant sneakily post-selecting the ensemble). Only then will the observed transition frequencies be taken as evidence for any law-like forwards transition probabilities.

Estimating backwards transition probabilities should proceed analogously. If we assume that the process has a limit distribution for $t \to -\infty$, then we cannot accept a sample as unbiased unless the initial distribution of the sample is in fact a limit distribution of the process. And if we assume that there is no limit distribution for $t \to -\infty$, then we are begging the question, because we have introduced a qualitative difference between forwards and backwards transition probabilities by hand.

Thus, while time-symmetric transition probabilities imply invariant equilibrium, a sample appropriate for estimating both forwards and backwards transition probabilities will be in equilibrium anyway. But now, the above criticisms all rely implicitly or explicitly on considering samples other than in equilibrium. Indeed, (A) uses convergence towards equilibrium (or the possibility of time-dependent distributions), so cannot be applied if the sample is in equilibrium already; (B) also requires the use of non-equilibrium ensembles because, trivially, forwards homogeneity and equilibrium imply backwards homogeneity; finally, (C) relies on considering alternative initial conditions, some of which will be non-equilibrium distributions.[5] The idea that convergence to equilibrium could be formally described using a time-symmetric stochastic process, plus a constraint on the initial distribution of the specific sample, is thus perfectly viable.

A case apart is provided by samples exhibiting what appear to be transient states. In the example, this is when we observe decay from the excited state to the ground state but no re-excitation, which is a case of particularly marked time-directed behaviour. At first sight, one might think that our argument above applies even in this case. Indeed, in order to have the forwards transition frequencies match the forwards transition probabilities, the sample must be totally decayed at the final time. By analogy, in order for the backwards transition frequencies to match the backwards transition probabilities, the sample must be totally decayed at the initial time (invariant distribution). But then, the samples exhibiting transience of the excited state are always biased for the purpose of estimating the backwards transition probabilities. There are two problems, however. Firstly, in a sample that is appropriate for estimating the transition probabilities in one direction of time, the transition frequencies in the opposite direction are partially ill-defined: thus, there are no samples appropriate for estimating both sets of transition probabilities (if such there be). Secondly and crucially, a non-zero initial frequency for excited states forces the backwards transition frequencies to be non-zero when the corresponding transition probabilities (assuming symmetry) should be zero, and thus is clearly not an allowable constraint.

A better way of treating samples with transient states will be to maintain that there is in fact a small but non-zero probability of re-excitation, which is a move analogous to standard reasoning in the deterministic case. (The fact that

[5]Conditionalising on two different equilibrium distributions (if there are several ergodic classes) will not yield different backwards transition frequencies, because the transition frequencies are fixed separately in each ergodic class.

Julius Caesar was alive and is now dead is not conclusive evidence against the time symmetry of classical mechanics.)

Recapitulating the above, we have seen that we can describe convergence to equilibrium using the transition probabilities of a stochastic process in equilibrium plus an assumption about special initial conditions (with an additional assumption in the case of apparently transient states). Therefore, the qualitative formal distinctions between forwards and backwards transition probabilities used as premises in the criticisms considered above are unwarranted.

We have not shown, however, that convergence to equilibrium can always be described using time-symmetric transition probabilities, because, other than in the two-state case, equilibrium is a necessary but not a sufficient condition for time symmetry. Indeed, there are also examples in which *circular currents* are called for: the transition matrices (2.28) above are stationary, so any initial distribution will converge to equilibrium, but in equilibrium there is a circular current. Intuitively, the 'atom' has a ground state 0 and two excited states 1 and 2, and state 2 decays to 0 directly with much larger probability than via the intermediate state 1. Thus, the transition probabilities fail to be time-symmetric.[6]

The import of these asymmetric cases can perhaps be minimised. The asymmetry appears to be more benign than in the criticisms considered above (e.g. if the forwards transition probabilities are time translation invariant, so are the backwards transition probabilities). Indeed, it does not appear that this asymmetry could justify a *qualitative* distinction between forwards and backwards transition probabilities. Furthermore, as briefly mentioned at the beginning of Section 2.3, the framework we have adopted allows us to describe these currents, but lacks any further structure that might explain them as determined perhaps by some underlying laws allowing a fuller analysis as regards time symmetry. Given such structure, the currents might turn out to be time-symmetric after all, in the sense that they would swap direction under time reversal of the underlying law.

A related example is provided by the inhomogeneous processes used in Nelson's (1966) approach to quantum mechanics. Without going into details, Nelson's approach is somewhat similar to the pilot-wave theory of de Broglie and Bohm, in that it takes quantum systems (in standard non-relativistic quantum mechanics) to be systems of point particles described in configuration space. Whereas de Broglie and Bohm take the velocity of the particles to be deterministically determined by the wave function of the system, Nelson postulates a stochastic process (a diffusion process) on the configuration space, and tries to impose conditions that would ensure that the process is determined in a certain way by the amplitude and phase of a complex function satisfying the Schrödinger equation. Whether or not Nelson's conditions achieve this, the process on configuration space definable through the wave function has as its current velocity the same velocity that arises in pilot-wave theory, which indeed changes sign with the time reversal of the Schrödinger equation. Thus, both time translation invariance and time symmetry, which are not apparent at

[6]My thanks to Iain Martel for making this point in conversation.

the level of the probabilities, are restored by the additional structure provided by the Schrödinger equation. Note that Nelson's approach can be adapted to the discrete case (Guerra and Marra, 1984). In this case the systematic component of the process is a probability current in the sense of (2.22), which again swaps sign under time reversal of the Schrödinger equation.[7]

While our above considerations apply only to processes that admit an invariant limit distribution, Nelson's processes generally have only an asymptotic distribution (also called equivariant), given by the usual quantum distribution $|\psi(\mathbf{x},t)|^2$ (similarly in Guerra and Marra's approach). We thus see that our considerations can be generalised to interesting cases of asymptotic convergence. That is, one can describe asymptotic convergence using a process that is time-symmetric – in the sense that the only time asymmetry is given by a current that swaps sign under time reversal – plus special initial conditions.[8]

2.4.3 Interpretation of Probability

We have tried to characterise the time symmetry of a Markov process in terms of forwards and backwards transition probabilities. To characterise similarly the *interpretation* of probabilities means that forwards and backwards transition probabilities would have the *same* or a *different* status. In particular, one could say that the idea of an (objectively) 'open future' and 'fixed past' means that forwards transition probabilities are law-like *chances*, while backwards transition probabilities are merely *epistemic*.

To say that both forwards and backwards transition probabilities are law-like seems less intuitive, since the two sets of probabilities determine the possible single-time distributions of the process (even uniquely), so the latter would also have to be taken as law-like. But law-likeness of probability distributions does not mean that relative frequencies have to always match the given probabilities. As long as an ensemble is finite, a law-like probability is compatible with infinitely many actual distributions, and it makes sense to consider *constraints* on, for instance, initial distributions or final distributions alongside with the laws. Indeed, the situation is quite analogous to that in the deterministic case. Deterministic laws determine the

[7] A more detailed introduction to Nelson's approach, including an explicit discussion of time symmetry and the status of the transition probabilities, is given in Bacciagaluppi (2005). As Nelson's approach relates to de Broglie and Bohm's pilot-wave theory, so Guerra and Marra's discrete case relates to the stochastic versions of pilot-wave theory, known as 'beable' theories, defined by Bell (1984).

[8] Observation in these cases, however, is definitely not classical. If one includes observers in the description (by adding some appropriate quantum mechanical interaction), when they gain knowledge about the state of the process, thus narrowing their epistemic distribution over the states, they effectively modify the wave function of the system, thus effectively modifying also the transition probabilities of the process, *both* forwards and backwards. (Note that convergence behaviour would thus be altered if monitored.)

time development of a system given, for instance, some initial condition; but which trajectory a system will actually follow is a contingent matter. Similarly, stochastic laws (whether symmetric or not) can be said to determine, in an appropriate sense, the time development of a system; but a stochastic process is a probability measure over a space of trajectories, and which trajectory the system will actually follow is a contingent matter. If we have a finite ensemble of systems, it is still a contingent matter which trajectories they will follow, regardless of whether the laws are deterministic or stochastic. (And, in fact, if the stochastic laws are assumed to be fundamental, then there is ultimately only one system – the universe – and only one trajectory.) Thus, at least as long as we are not dealing with literally infinite ensembles, we can make the same distinction between law-like time development and contingent initial or final states, or distributions over states, in the case of both deterministic and stochastic laws, and this even if we assume that both forwards and backwards transition probabilities are law-like, despite the ensuing law-likeness of single-time distributions.[9]

We can imagine a stochastic world in which observed transition frequencies typically show not merely a quantitative but a qualitative difference between forwards and backwards transition frequencies, as in the examples in Section 2.4.1. However, our analysis in Section 2.4.2 shows that arguments from observed frequencies fail to establish an asymmetry between the corresponding probabilities: although ensembles that are not in equilibrium lead to distorted frequencies, neither the preponderance of non-equilibrium ensembles in such a world nor any conclusions drawn on the basis of these frequencies can be arguments against time-symmetric transition chances (and this despite the fact that equilibrium is a necessary condition for (2.19)). The only serious source of time asymmetry at the level of the formalism and therefore potential motivation for a time-asymmetric interpretation would seem to be the presence in some cases of circular currents, which indeed yield quantitatively asymmetric transition probabilities. However, circular currents yield no *qualitative* difference that could justify a different status for forwards and backwards transition probabilities. In particular, if the only difference between past and future is the presence of a current in one direction or another along a closed chain of states, it is difficult to see which of the two directions should correspond to an *open* ‘future’ as opposed to a *fixed* ‘past’. Thus, the possibility of an asymmetry in terms of circular currents does not seem to be of the kind that would justify a time-asymmetric interpretation of probability.

---

[9]The notion of a constraint is of course more intuitive when one is talking about a subsystem on which one performs experiments (as in thermodynamics or statistical mecha nics when compressing a gas into a small volume), but it is meant to apply generally. As emphasised by the anonymous referee, in the case of a stochastic theory such constraints will not only be ‘special’ in some sense but they will be improbable in the sense specified by the process itself. The further question of whether and how the contingent trajectories (or distributions) should be explained thus acquires a new twist as compared to the deterministic case.

At least in the case of processes with an invariant limit distribution, our analysis suggests that both forwards and backwards transition probabilities can be considered law-like. Therefore, whatever approach to the foundations of probabilities one might take, a *time-symmetric interpretation of probabilities* appears to be a natural option in the context of classical Markov processes.

Acknowledgements The first version of this paper was written while I was an Alexander-von-Humboldt Fellow at the Institut für Grenzgebiete der Psychologie und Psychohygiene (IGPP), Freiburg i. Br. I wish to thank in particular Werner Ehm at IGPP, Iain Martel, then at the University of Konstanz, and David Miller at the Centre for Time for useful discussions and suggestions, as well as Mauricio Suárez for the kind invitation to contribute to this volume and an anonymous referee for interesting comments.

2.1 Appendix

We now prove that symmetry of the transition probabilities (2.19), together with the further assumptions that the state space is finite and that the transition probabilities are continuous, implies equilibrium of the process.

We proceed by induction on the size n of the state space. The case $n = 1$ is trivial. Assume that the result has been proved for all sizes $1 \leq m < n$. We now prove it for n by *reductio*.

Assume that the single-time distribution is not invariant, i.e.

$$\exists s \exists t \geq s, \quad \mathbf{p}(t) = P(t|s)\mathbf{p}(s) \neq \mathbf{p}(s) \; . \tag{2.51}$$

For the rest of the proof we now fix such an s.

Since we assume (2.19), i.e. $P(t|s) = P(s|t)$, we also have

$$\mathbf{p}(s) = P(t|s)\mathbf{p}(t) \; , \tag{2.52}$$

and therefore

$$P(t|s)^2\mathbf{p}(s) = \mathbf{p}(s) \qquad \text{and} \qquad P(t|s)^2\mathbf{p}(t) = \mathbf{p}(t) \; . \tag{2.53}$$

Now fix a time $t \geq s$ and consider the matrix $P := P(t|s)^2$. This is an $n \times n$ stochastic matrix that we can consider as the transition matrix of a homogeneous Markov process with discrete time. By (2.53), $\mathbf{p}(t)$ and $\mathbf{p}(s)$ are both invariant distributions for this Markov process, and by (2.51) they are different.

By the ergodic theorem for discrete-time Markov processes, existence of at least two different invariant distributions implies that there are at least two ergodic classes. Therefore (whether or not there are any transient states), P must have a block diagonal form

$$P = \begin{pmatrix} P' & \mathbf{0} \\ \mathbf{0} & P'' \end{pmatrix} \; , \tag{2.54}$$

where P' is an $m \times m$ matrix and P'' an $(n-m)\times(n-m)$ matrix, for some $0 < m < n$.

For fixed s, $P = P(t|s)^2$ depends on t, and so *a priori* could m; but in fact $m(t)$ is independent of t. Indeed, assume there is an $m \neq m(t)$ such that for all $\varepsilon > 0$ there is a t' with $|t - t'| < \varepsilon$ and $m(t') = m$. The matrix elements of $P = P(t|s)^2$, in particular the ones off the diagonal blocks, are continuous functions of the transition probabilities. Therefore, by the continuity of the transition probabilities, $P(t|s)^2$ must also have zeros off the same diagonal blocks, i.e. $m = m(t)$, contrary to assumption. Therefore, for each $m \neq m(t)$ there is an $\varepsilon(m) > 0$ such that for all t' with $|t - t'| < \varepsilon(m)$ we have $m(t') \neq m$. Taking the smallest of these finitely many $\varepsilon(m) > 0$, call it ε_0, it follows that $m(t') = m(t)$ for all t' in the open ε_0-neighbourhood around t. However, again by the continuity of the matrix elements, this ε_0-neighbourhood is also closed, and therefore it is the entire real line. Since t was arbitrary, $P(t|s)^2$ has the form (2.54) with the same m for all $t \geq s$.

We now focus on the matrix $P(t|s)$ itself rather than on $P(t|s)^2$. Assume that for some $t \geq s$ it has some element $p_{k|l}(t|s)$ outside of the $m \times m$ and $(n-m)\times(n-m)$ diagonal blocks. In order for $P(t|s)^2$ to have the given block diagonal form, several other elements of $P(t|s)$ have to be zero, in particular all elements in the k-th column of $P(t|s)$ that lie inside the corresponding diagonal block.

Since $P(t|s)$ is a stochastic matrix and every column sums to 1, it follows that already those elements of the k-th column that lie outside the diagonal blocks sum to 1, and therefore the sum of all elements in the diagonal blocks of $P(t|s)$, call it $d(t)$, is at most $n - 1$, i.e.

$$d(t) = \sum_{i,j \leq m} p_{i|j}(t|s) + \sum_{i,j \geq m+1} p_{i|j}(t|s) \leq n - 1 \,, \tag{2.55}$$

for any $t \geq s$ such that $P(t|s)$ has some element outside of the diagonal blocks. Let t_0 be the infimum of such t. By continuity, we have also

$$d(t_0) \leq n - 1 \,. \tag{2.56}$$

Now, if $t_0 \neq s$, then for all $t < t_0$ we have that $d(t) = n$, but then by continuity $d(t_0) = n$, contradicting (2.56). If instead $t_0 = s$, since $P(s|s) = \mathbf{1}$, we again have $d(t_0) = n$, contradicting (2.56). For all $t \geq s$, thus, $P(t|s)$ has the same block diagonal form as $P(t|s)^2$ with fixed m.

But then, our original Markov process decomposes into two sub-processes, with state spaces of size m and $n - m$, respectively. If $\mathbf{p}(t) \neq \mathbf{p}(s)$ (assumption (2.51)), then the same must be true for at least one of the two sub-processes, but, by the inductive assumption, this is impossible. Therefore, (2.51) is false and

$$\forall s \forall t \geq s, \; \mathbf{p}(t) = \mathbf{p}(s) \,, \tag{2.57}$$

QED.

References

Arntzenius, F. (1995), Indeterminism and the direction of time, *Topoi* 14, 67–81.

Bacciagaluppi, G. (2005), A conceptual introduction to Nelson's mechanics, In R. Buccheri, M. Saniga and E. Avshalom (eds.), *Endophysics, Time, Quantum and the Subjective*, Singapore: World Scientific, pp. 367–388.

Bacciagaluppi, G. (2007), Probability, arrow of time and decoherence, *Studies in History and Philosophy of Modern Physics* 38, 439–456.

Bell, J. S. (1984), Beables for quantum field theory, CERN-TH. 4035/84. Reprinted in J. S. Bell, *Speakable and Unspeakable in Quantum Mechanics*, Cambridge: Cambridge University Press, pp. 173–180.

Doob, J. L. (1953), *Stochastic Processes*. New York, NY: Wiley.

Guerra, F. and Marra, R. (1984), Discrete stochastic variational principles and quantum mechanics, *Physical Review* D 29, 1647–1655.

Hajnal, J. (1958), Weak ergodicity in non-homogeneous Markov chains, *Proceedings of the Cambridge Philosophical Society* 54, 233–246.

Kolmogorov, A. (1931), Über die analytischen Methoden in der Wahrscheinlichkeitsrechnung, *Mathematische Annalen* 104, 415–458.

Nelson, E. (1966), Derivation of the Schrödinger equation from Newtonian mechanics, *Physical Review* 150, 1079–1085.

Sober, E. (1993), Temporally oriented laws, *Synthese* 94, 171–189.

Uffink, J. (2007), Compendium of the foundations of classical statistical physics, In J. Butterfield and J. Earman (eds.), *Handbook of the Philosophy of Physics*, Part B Amsterdam: North-Holland, pp. 923–1074.

Watanabe, S. (1965), Conditional probability in physics, *Progress of Theoretical Physics Supplement*, Extra issue (1965), 135–167.