7

Humean Chance in Physics

Coauthored with Roman Frigg[*]

In Carl Hoefer: *Chance in the World. A Humean Guide to Objective Chance*. New York: Oxford University Press, 2019, 181-213.

7.1. Introduction

After the acceptance of Newton's mechanics and the development of its increasingly sophisticated mathematical formulations, it was for over a century common to assume that determinism reigns in the physical world. Probability was no part of mainstream fundamental physical theory. Nevertheless, probability *was* introduced into physics during the height of the Newtonian paradigm, and its applications in physics—both at the level of theory, and in experimentation and instrumentation—have grown since by leaps and bounds to the point where it is hardly an exaggeration to say that probability is near-ubiquitous in modern physics.

In our view, the Humean approach to objective probability defended in this book is well suited to explicating the uses of objective probability that we have encountered in both theoretical physics (from quantum mechanics and statistical mechanics to modern stochastic-dynamics theories) and experimental physics (the various phases of experiment and equipment design, such as calibration and shielding techniques, data analysis, e.g., noise modeling, and so on). It is also well suited to understanding and justifying Monte Carlo techniques, which are used in many areas of physics practice. There is no way that we can survey all these applications in one chapter, however, and we will not make the attempt. Instead, we begin by observing that there are two basic ways in which objective probabilities seem to occur in physical theories: they can occur at a fundamental or "rock-bottom" level in the theory, or they can be *superimposed* on a deterministic fundamental

[*] *Department of Philosophy, Logic and Scientific Method and Centre for Philosophy of Natural and Social Sciences (CPNSS), London School of Economics*

level. In this chapter we will discuss paradigmatic examples of these two ways, starting with Boltzmannian statistical mechanics (SM), and then moving on to the standard interpretation of quantum mechanics (QM). These two cases, we believe, provide the template for understanding probabilities in other parts of physics that we cannot discuss here.

When probability was introduced in a quite crucial way into the physics of atomic and molecular systems, in Maxwell's and Boltzmann's works on statistical mechanics, it was possible to see its role as still essentially epistemic and non-fundamental, and this remains a viable view concerning classical statistical mechanics today (although not one we endorse). The rise of quantum mechanics in the 1920s famously changed this situation and turned things upside down: suddenly the *fundamental* physics of the micro-world was seen as genuinely indeterministic, governed by laws permitting only probabilistic predictions. Objective chances, it seemed, were both needed in *and provided by* fundamental physics; and it seemed that these chances, moreover, *could not be understood as merely epistemic*, in any of the ways familiar from nineteenth-century physics.

As we will see in this chapter, all the claims concerning probabilities in the preceding paragraph, for both SM and QM, are highly controversial and are still actively disputed in both the physics and philosophy of physics communities. But two facts are beyond dispute, and these facts make it highly desirable that we explore whether, and if so how, HOC can meet the challenge of delivering the objective probabilities of QM and SM. The first fact is this: both SM and QM (and later quantum field theories) have been enormously successful empirically and fruitful theoretically, and the probabilities they introduce are indispensable for these successes. The probabilities also generate hypotheses that are put to experimental test, and these hypotheses have been found to be in exact agreement with empirical results time and again. A natural explanation of this empirical success is that these probabilities capture "how things are," rather than just offering a codification of our own ignorance.

The second fact is that in both theories the probabilities are set in an "a priori" fashion, that is to say, derived from the theory's basic equations and models, not derived inductively from observation of frequencies or patterns of frequencies. This "*quasi*-a-priori"

status is *prima facie* a challenge for a Humean account of all objective probabilities. For if objective chances are, so to speak, *written directly into the laws of nature*, this seems to challenge both the skeptical arguments of chapter 1, and the very need for a Humean/reductive account of probability (unless one is already committed to reductive Humeanism concerning the laws of nature in general, as is the case for some philosophers, but not for us). And even if SM probabilities are not taken to be directly written into the laws of nature, it is still desirable to show that the Best System recipe is likely to yield precisely the objective probabilities of SM (neither leaving them out of the picture entirely, nor amending them very significantly). In this chapter we will confront these challenges by exploring the ways in which HOC can be seen to capture both SM and QM probabilities.

7.2. Classical Statistical Mechanics

In this section we discuss the thorny problem of how to understand the probabilities introduced into physics by classical SM. Traditional approaches to probability, such as propensity views and frequentist accounts, run into serious difficulties, nor do special reductive definitions found in the theory (e.g., time averages) work at all well.[1] But as we will see, the pragmatic Humean approach seems to square nicely with the ontology and ideology of SM.

7.2.1. Background of Classical SM

The behavior of macroscopic systems such as a gas in a box is, to a very high degree of accuracy, correctly described by thermodynamics (TD). TD characterizes the states of systems in terms of variables like temperature, pressure, and volume, which pertain to the system as a whole, and it posits that processes have to be such that the entropy of

[1] Against a propensity interpretation of SM probabilities, see (Clark, 2001). Against the time averages approach, see (von Plato, 1981, 1982) and (van Lith, 2001).

an isolated system does not decrease. At the same time, the system can be regarded as a collection of molecules, each governed by the laws of mechanics, which in what follows we will (unless otherwise noted) assume to be the laws of classical mechanics (CM). Statistical mechanics aims to establish a connection between these two ways of looking at the system and to account for TD behavior in terms of the dynamical laws governing the microscopic constituents of macroscopic systems and probabilistic assumptions. We now briefly review classical SM and discuss in some detail how probabilities are introduced into that theory. SM comes in different versions and formulations. We here focus on what has become known as Boltzmannian SM and refer the reader to (Frigg, 2008), (Sklar, 1993), and (Uffink, 2006) for discussions of other approaches.

To introduce SM we first have to review briefly the main tenets of CM. CM describes the world as consisting of point-particles, which are located at a particular point in space and have a particular momentum (where a particle's momentum essentially is its velocity times its mass). A system's state is fully determined by a specification of each particle's position and momentum—that is, if you know the positions and the momenta of all particles in the system, you know everything that there is to know about the system's state from a mechanical point of view. Conjoining the space and momentum dimension of all particles of a system in one vector space yields the so-called phase space Γ of the system. For a particle moving around in the three-dimensional space of our everyday experience, the phase space basically consists of all points $X = (x, y, z, p_x, p_y, p_z)$, where x, y, and z are the three directions in space, and p_x, p_y, and p_z are the momenta in these directions. So the phase space of one particle has six (mathematical) dimensions. The phase space of a system consisting of two particles is the collection of all points $X = (x_1, y_1, z_1, x_2, y_2, z_2, p_{x_1}, p_{y_1}, p_{z_1}, p_{x_2}, p_{y_2}, p_{z_2})$, where x_1, y_1, and z_1 are the spatial locations of the first particle, x_2, y_2, and z_2 the one of the second particle, and $p_{x_1}, \ldots,$ are the momenta in the respective directions. Hence, the phase space of such a system is 12-dimensional. The generalization of this to a system of n particles—which is what SM studies—is now straightforward: it is a $6n$-dimensional abstract mathematical space. If X is the state of an n particle gas, it is also referred to as the system's *micro-state*. So, the micro-state of a system

consisting of n particles is specified by a point x in its $6n$-dimensional phase space Γ.

An important feature of Γ is that it is endowed with a so-called Lebesgue measure μ. Although Γ is an abstract mathematical space, the leading idea of a measure is exactly the same as that of a volume in the three-dimensional space of everyday experience: it is a device to attribute sizes to parts of space. We say that a certain collection of points of this space (for instance, the ones that lie inside a bottle) have a certain volume (for instance, one liter), and in the same way can we say that a certain set of points in Γ has a certain μ-measure. If A is a measurable set of points in Γ, we write $\mu(A)$ to denote the μ-measure of this set. At first it may seem counterintuitive to have measures ("volumes") in spaces of more than three dimensions (as the preceding X shows, the space of a one-particle system has 6 and that of a two-particle system 12 dimensions). However, the idea of a higher dimensional measure becomes rather natural when we recall that the moves we make when introducing higher dimensional measures are the same as when we generalize one-dimensional length, which *is* the Lebesgue measure in one dimension, to two dimensions, where the surface area is the Lebesgue measure, and then to three dimensions, where volume is the Lebesgue measure.

The state of a system will generally change over time in the way prescribed by classical mechanics. Since each particle moves along a continuous line in space, the system's micro-state in phase space similarly moves along some continuous path in phase space. The "line" that $\phi_t(X)$ traces through the phase space is called a trajectory. What trajectory a system follows depends on where the system starts. The state X where the system begins its motion at time t_0 (the moment when the process begins) is called the "initial condition." The system's trajectory is governed by the so-called *Hamiltonian equations of motion*. CM thus determines how a system located at any point in the phase space will evolve—move to a new point—as a function of time. The function that tells us what the system's state at some later point will be is called a "phase flow" and we denote it with the letter ϕ. We write $\phi_t(X)$ to denote the state into which X evolves under the dynamics of the system if time t (e.g., one hour) elapses, and similarly we write $\phi_t(A)$ to denote the image of a set A (of states) under the dynamics of the system.

More precisely, Hamilton's equations of motion define a *measure-preserving* phase flow ϕ_t on Γ. That is, $\phi_t : \Gamma \rightarrow \Gamma$ is a one-to-one mapping for every real number t and $\mu(\phi_t(R)) = \mu(R)$ for every measurable region $R \subseteq \Gamma$. In what follows, we assume that the relevant physical process begins at a particular instant t_0 and we adopt the convention that $\phi_t(x)$ denotes the state of the system at time $t_0 + t$ if it was in state x at t_0, and likewise $\phi_t(R)$; x is then commonly referred to as the "initial condition."

In (Boltzmannian) SM a key notion is that of a "macro-state." One assumes that every system has a certain number of macro-states $M_1, \ldots, M_k$ (where k is a natural number that depends on the specifics of the system), which are characterized by the values of macroscopic variables; in the case of a gas these would be pressure, temperature, and volume. It is a basic posit of Boltzmannian SM that the M_k supervene on the system's micro-states. Therefore each macro-state M_k is associated with a macro-region $\Gamma_k \subseteq \Gamma$, so that the system is in macro-state M_k at t *iff* its micro-state x lies within Γ_k at t.[2] For reasons that will become clear soon, we choose special labels for the first and the last macro-state: $M_1 = M_p$ and $M_k = M_{eq}$

In Figure 7.1 we see a schematic representation of the phase space of a system consisting of gas molecules in a box. Initially the system's molecules are all confined to the left side of the box. At time t_0 the partition is removed and the gas expands to fill the whole box; that is, its micro-state evolves from someplace inside M_p to end up somewhere in M_{eq}, the equilibrium macro-state. This evolution of the system takes it from a low-entropy initial state (in the instant after the partition is removed) to a high-entropy state. Showing that such an entropy-increasing evolution is overwhelmingly likely is a key aim of SM.

In TD, entropy is a notion defined in terms of macroscopic features of a system, so it makes sense that in SM entropy is defined for our

[2] Although thermodynamic properties such as temperature and pressure are continuous rather than discrete quantities, it is useful in SM to treat them in a "coarse-grained" fashion, so that a specific integer k indicates systems having temperature within a certain small range of values, pressure within a certain small range of values, etc. This gives the macro-regions of phase space corresponding to the macro-states M_k non-trivial volume, and is useful in other respects as well. See (Frigg, 2008) for details.

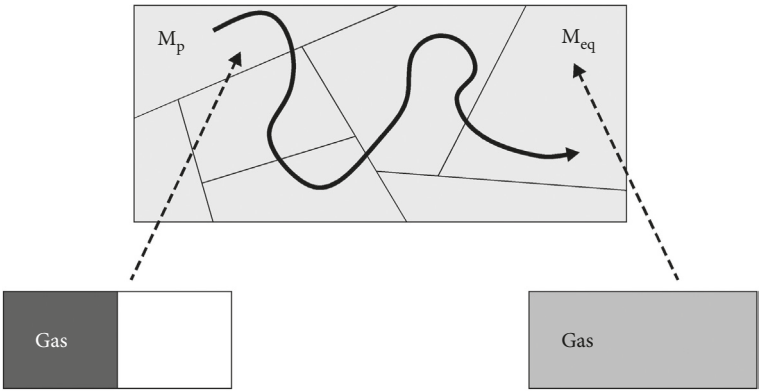

Figure 7.1. A micro-state starting in M_p, far from equilibrium, evolves through various intermediate macro-states and finally reaches M_{eq}, the equilibrium macro-state.

macro-states M_k. The Boltzmann entropy of a macro-state is defined as $S_B(M_k) = k_B \log(\mu(\Gamma_k))$, where k_B is the Boltzmann constant. Since the system is only exactly in one macro-state/macro-region at any moment, we can define the Boltzmann entropy of the *system* at time t simply as the entropy of the macro-state at time t: $S_B(t) = k_B \log(\mu(\Gamma_t))$, where Γ_t is the macro-region in which the system's micro-state is located at time t. It can be shown that, unlike the depiction in Figure 7.1, the equilibrium state M_{eq} is *by far* the largest of all states (under μ), a fact known as the dominance of the equilibrium state. In fact, for large n it is vastly larger than the area of all other regions.[3] Since the logarithm is a monotonic function, it follows that the Boltzmann entropy is maximal for the equilibrium macro-state.

The Second Law of TD says, very roughly, that the TD entropy of a closed system cannot decrease. In SM it is common to give a somewhat stronger reading to this law and take it to assert that a system, when left to its own, approaches equilibrium (and hence its entropy

[3] The notion of being large can be explained in different ways and different justifications can be given. For a discussion of this point see (Werndl & Frigg, 2015).

approaches its maximum value). Explaining this approach to equilibrium, then, is the aim of SM.

However, this aim is only of any interest if the system originally starts off in a non-equilibrium state. That this be the case is the subject matter of the so-called *Past Hypothesis*, the postulate that the system came into being in a low entropy state, which we call the *past state*. (The term "Past Hypothesis" is sometimes reserved for approaches in which the system under consideration is the entire universe, in which case it says that the universe came into being in a low-entropy macro-state, which is provided to us by modern Big Bang cosmology. We return later to the issue of the nature of systems studied in SM. In the meantime we use "Past Hypothesis" as referring to the initial state of a system, irrespective of the precise nature of the system.) Since the past state and the equilibrium state are of particular importance, we introduce special labels and denote the former by M_p (which is associated with Γ_p) and the latter by M_{eq} (which is associated with Γ_{eq}).

7.2.2. Explaining the Approach to Equilibrium

There are two different schools of thought that have divergent understandings of what exactly an explanation of the approach to equilibrium amounts to, and, accordingly, propose different solutions. One approach, the *TD-likeness approach*, will be explained in the following. The other, which we might call the *transition probabilities approach*, is associated with the work of Albert (2000). For the sake of brevity we concentrate on the first approach; readers are referred to (Frigg 2008) and (Frigg & Hoefer, 2015) for our take on the second approach.

The TD-likeness approach has in recent days been advocated by Lavis (2005) and can be traced back to Boltzmann himself. According to this approach, a justification of the second law amounts to showing that the system is highly likely to exhibit thermodynamic-like behavior (TD-like behavior). We have TD-like behavior *iff* the entropy of a system that is initially prepared in a low-entropy state increases until

it comes close to its maximum value, and then stays there and only exhibits frequent small and rare large fluctuations away from equilibrium (as depicted in Figure 7.2).

Before elucidating the notion of "highly likely" we have to introduce the concept of ergodic motion. Roughly speaking, the motion of a system is ergodic if for any set $A \subseteq \Gamma$, the proportion of time the trajectory spends in A is equal to the proportion of the measure A takes up in Γ in the long run (for instance, if A occupies a quarter of Γ, then the system spends a quarter of the time in A). If a trajectory is ergodic, then the system behaves TD-like because the dynamics will carry the system's state into Γ_{eq} and will keep it there most of the time. The system will move out of the equilibrium region every now and then and visit non-equilibrium states. Yet since these are small compared to Γ_{eq} it will only spend a small fraction of time there. Accordingly, the entropy is close to its maximum most of the time and fluctuates away from it only occasionally.

As a matter of fact, whether or not a system's motion is ergodic depends on the initial condition: some initial conditions lie on trajectories that are ergodic, while others don't. This realization is the clue to introducing probabilities. Consider an arbitrary subset $C \subseteq \Gamma_p$. We may postulate that the probability that the initial condition X lies within C at time t_0 is

$$p(C) = \frac{\mu(C)}{\mu(\Gamma_p)} \qquad \text{(Eq. 2)}$$

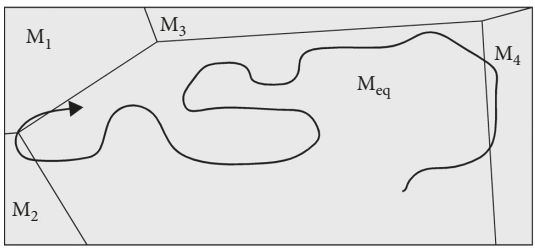

Figure 7.2. Time evolution of a system in equilibrium. It remains in M_{eq}, aside from brief sojourns into nearby non-equilibrium macro-states.

Let us refer to this principle as the Past Hypothesis Proportionality Postulate (PHPP). If we now denote by $E \subseteq \Gamma_p$ the subset of all initial conditions that lie on ergodic trajectories, then $p(E)$ is the probability that the system behaves TD-like:

$$p\big(\text{TD-like behavior}\big) = p(E) \qquad \text{(Eq. 3)}$$

If the values of $p(E)$ come out high, then one has justified the Second Law. Whether or not this is the case in actual systems is a non-trivial and controversial question. For the sake of argument let us assume that it is.[4]

7.2.3. Humean Chances in SM

We now argue that SM probabilities can be interpreted as HOCs. We begin to see how such chance and determinism are compatible if we pay attention to the curious hybrid character of Equation 3. The term on the right-hand side comes from the fundamental micro-theory, namely classical mechanics. The term on the left-hand side refers to probabilities for macroscopically observable events; in fact, we can sometimes see with our unaided eyes whether a gas in a container behaves thermodynamically. Given this, why do we need micro-physics at all to attribute probabilities to these events? We assign probabilities to everyday events like getting heads when flipping a coin, seeing the sun shine tomorrow, cracking the jackpot in Sunday's lottery, seeing our favorite team win the next game, etc., without ever making reference to an underlying micro-theory.

Take the example of a coin toss, discussed at length in chapter 3. The event-type we call "a good flip of a fair coin" is widespread in HM around here. Furthermore, it is a fact, first, that in HM the relative frequency of each discernible side-type landing upward is very close to 0.5 and, second, that there are no easily discerned patterns

[4] In a recent paper, Roman Frigg and Charlotte Werndl (2011) argue that it is.

to the flip outcomes (it is not the case, for instance, that a long sequence of outcomes consist of alternating heads and tails). Does a Best System for our HM contain a rule assigning probabilities to coin-flip outcomes, and if so, what probabilities will the rule assign? The second question is easy to answer: the rule that postulates $p(H) = p(T) = 0.5$ is the one that has best fit and simplicity. But does such a rule belong in the system at all? *Prima facie* the answer should be "yes," since adding such a rule extends the system to a new and widespread type of phenomenon in HM, at a tiny cost in reduced simplicity. As we noted in chapter 3, there may be an even better chance rule that could be part of the Best System, which would embrace coins and dice and tetrahedra and dodecahedra and other such symmetric, flippable/rollable solids. The rule would say that where such-and-such symmetry is to be found in a solid object of middling size with n possible faces that can land upward (or downward, thinking of tetrahedra), and when such objects are thrown/rolled, the chance of each distinct face being the one that lands up (or down) is exactly $1/n$. Given what we know about dice and tetrahedra and so forth, it is quite plausible that this rule belongs in the Best System; and it entails the coin-flip chances. So it enhances both simplicity and strength without much loss in fit, and hence on balance it is better to add this rule to the system rather than a set of rules, one for each type of regular n-sided solid. Hence, the chance of heads on a fair flip of a coin would seem certainly to exist, and be 0.5, in a Best System for our world.

The same kind of reasoning applies to gases. Behaving TD-like is a macro-property in much the same way as showing heads, and so we can introduce chances for that type of event in the same way. Preparing a gas in M_p and then letting it evolve unperturbed corresponds to flipping a coin; the gas behaving TD-like or not correponds to getting heads or tails. HM contains many gases that were first prepared in M_p and then left to themselves and so we have solid frequency data to rely on (just as in the case of coins). The overwhelming majority of those behave TD-like and so we formulate the rule p(TD-like) = 0.9999 and p(non-TD-like) = 0.0001 (omitting many 9s and 0s here for brevity). Let us refer to these as the *macro-probability rules* for TD-like and non-TD-like behavior, respectively. This rule is simple and has good fit. It also possesses strength because it turns out that it holds true not

only for gases; in fact, liquids and solids also behave TD-like most of the time and so the rule is applicable universally (or almost universally at any rate). For this reason our rule is part of the Best System, and the probabilities it introduces are chances in the sense of HOC.

But notice that something curious has happened. We started explaining the hybrid character of Equation 3, and ended up making claims about the probabilities that appear in Equation 3 without reference to that equation at all! Equation 3 *seems* to have become an idle wheel. Has it really? The answer is "yes" and "no."

"Yes" because, and this is an important point, probabilities for macro-events like coin flips and the behavior of gases can (a) be defined without reference to the underlying micro-physics and (b) be genuine parts of the Best System. Especially the second part of this claim may strike some as unfounded, and we will come back to it later. Let us for now accept that macro-chances are genuine chances. The "no" part of the preceding answer holds that Equation 3 is not an idle wheel *even if* one accepts that macro-chances are real and genuine chances.

How can that be? The crucial point to realize is that, first appearances notwithstanding, Equation 3 does not *give* us the chance for TD-like behavior. We don't need to be given anything—we have the chance (via the *macro-probability rule*). Rather, Equation 3 is both a consistency check and an explanation. We don't want to place too much emphasis on the latter and mainly focus on the former, but there is the pervasive intuition that if a macro-result can be derived from a more fundamental theory, there is explanation. Those who share this intuition—among them us—can see Equation 3 as providing an explanation of the *macro-probability rule* for TD-like behavior (because set *E*, which is the key ingredient of that equation, is in fact given to us by the dynamical hypothesis that the system be ergodic).[5]

[5] The character and the strength of the explanation provided depends, of course, on how one views several things, including: the fundamentality of SM; the status of fundamental mechanical laws (whether of CM or QM); and the success of ergodicity-based explanations of thermodynamic behavior. For an account of the latter item, see (Frigg & Werndl 2011).

Let us now turn to consistency. The different parts of a Best System have to be consistent with each other. For this reason, whenever a macro-level chance rule and a micro-level chance rule are extensionally equivalent, then the chances they ascribe must agree, or be very nearly in agreement. This, of course, does not rule out the possibility of minor adjustment. For example, assume we adopted the 50/50 rule for heads and tails when flipping a coin. Now suppose we know for sure that we get the reductive relations right and we have the correct micro-theory, and based on these we find 49/51. This is no real conflict because there is some flexibility about the macro-chances, and if there are very good overall reasons for making adjustments, then the Humean can make these. But there is a breaking point: if the micro-theory predicts 80/20, we have to go back to the drawing board.

Recall Equation 3: p (TD-like behavior) = $p(E)$. This equation provides a consistency check in two ways. The first and more straightforward one is that a calculation has to show that $p(E) = \mu(E) / \mu(\Gamma_p) = 0.9999$, or at least that the value of $p(E)$ is very close to 0.9999. If the system is ergodic we find that $p(E) = 1$ because initial conditions that lie on trajectories for which time and space averages do not coincide form a set of measure zero. So this requirement is met. The second way has to do with the supervenience of the measure in Equation 3. It is one of the posits of HOC that chance functions supervene on the Humean Mosaic. But before we can see what consistency constraint emerges from that posit, more needs to be said about what it means for a rule like Equation 3 to supervene on the HM.

To understand what it means for Equation 3 to supervene on the HM, we first have to ask the same question about Equation 2—i.e., $p(C) = \mu(C) / \mu(\Gamma_p)$. To begin with, notice that this equation in effect expresses the conditional probability of finding a system's microstate in set C given that it lies in Γ_p (because by assumption $C \subseteq \Gamma_p$). Now look at the same conditional probability from the point of view of HOC. There is a well-circumscribed class of objects to which a chance rule like PHPP applies (gases, etc.). Each of these, we are assuming here, is a classical system with a precise initial condition X at t_0, which, by assumption, lies within Γ_p. Now go through the entire HM and put every single initial condition x into Γ_p. The result of this

is a swarm of points in Γ_p. Then recall that HOC can be viewed as a sophistication of finite frequentism, and chances should closely track relative frequencies wherever such frequencies are available in large quantity. Hence the chance of an initial condition being in set C given that it lies in Γ_p should be close to the fraction of points in set Γ_p that lie in C (in the same way in which the chance of heads is the fraction of heads in the set of all coin toss outcomes).

But listing all points individually and checking whether they lie in C is extremely cumbersome and won't make for a simple system. So we have to reduce the complexity of the sytem by giving a simple summary of the distribution of points. To this aim we approximate the swarm of points with a continuous distribution (which can be done using one of the well-known statistical techniques) and normalize it. The result of this is a probability density fuction ρ on Γ_p, which can be regarded as an expression of the "initial condition density" in different subsets C of Γ_p. This distribution supervenes

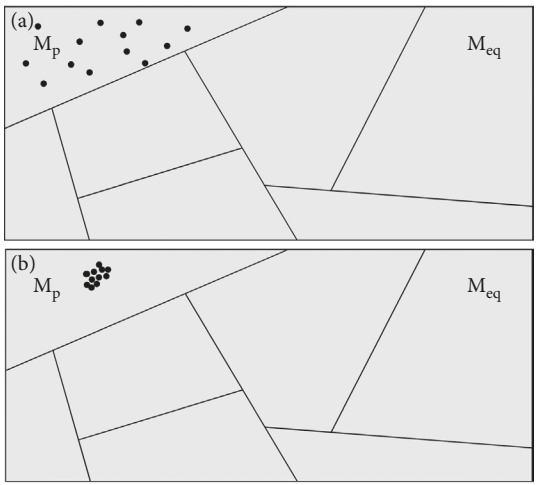

Figure 7.3*a,b*. Actual initial conditions of systems that start out in M_p. If evenly distributed (*a*), the Lebesgue measure will provide a good probability measure, as in Equation 2. If the actual initial conditions are too clustered (*b*) this will not generally be the case.

on the HM by construction. The consistency constraint now is that $\rho(C)$ be equal to (or in very close agreemet with) $\mu(C) / \mu(\Gamma_p)$ for all subsets C of Γ_p. This is a non-trivial constraint. For it to be true, it has to be the case that the initial conditions are more or less evenly distributed over Γ_p because $\mu(C) / \mu(\Gamma_p)$ is a flat distribution over Γ_p. This is illustrated in Figure 7.3 *a*. If it turned out that all initial conditions were crammed into one corner of Γ_p, as in Figure 7.3 *b*, then Equation 2 would have poor fit, which would preclude its being part of the Best System, despite its great simplicity. So the requirement $\rho(C) = \mu(C) / \mu(\Gamma_p)$ presents a real touchstone for the system. Equation 3 is then dealt with easily. Since Equation 2 has to hold for all C, a fortiori it has to hold for E. Hence Equation 3 gives the correct chance for TD-like behavior.

This leaves us with the question of why the Best System for our world might have macro-chances for TD-like behavior, if the SM chances emerge as part of the system in the way sketched in the preceding. Would they not be superfluous?

We think not. Even if the world is classical at bottom and classical mechanics is the fundamental theory of the universe, it does not follow that everything that can be said about HM has to be said in the language of the fundamental theory. As argued in chapters 3–5, probability rules can be formulated in a macro-language pertaining to a certain level of discourse, and probability rules thus introduced have equal right to be considered for inclusion in a Best System package of rules, alongside micro-level rules. On our view, then, the $1/n$ rule for gambling devices and the *macro-probability rule* for TD-like behavior are genuine chance rules because they improve the score, in terms of the three basic metatheoretical virtues, of a system that includes them, compared to a system that does not and (say) only contains micro-level chance rules.

The key point is that our pragmatic Best System account characterizes *simplicity* in a way that connects with the way real sciences work. A system that has rules at various levels and for various domains, not just micro-physical-level rules, is one that has more relevance to real science. It is a matter of fact that different scientific disciplines talk in their own language, and we have no reason to believe that this will change. Biology speaks its own language, and there

is no indication that progress in biology would mean that biologists start formulating their theories in the terms of fundamental physics. A view of chance (or laws of nature, for those who adopt a HBS view of laws too) that ignores this aspect of science is one that is of little use to science.[6]

One might worry that adding probability rules in a macro-language to the system in fact renders the system less simple without adding any strength, at least if the system already has micro-level rules that extensionally "cover" all the phenomena covered by the macro-level rule (as one might say when comparing Equation 3 with the macro-probability rule introduced earlier). This is mistaken. In fact, adding these macro-level rules can make the system simpler! The reason for this is that aspect of simplicity mentioned in earlier chapters, simplicity in derivation: it is hugely costly to start from first principles every time you want to make a prediction about the behavior of a macro-object. So the system becomes simpler in that sense if we write in rules about macro-objects.

Some philosophers would object that the chance for macro-level event types cannot be independent of the micro-physics of the world. Surely, the argument goes, there must be some dependence there! If the physics of our world were vastly different from what it is, then the chance for coin flips to land heads would be different too. So one cannot just introduce chances at a macro-level and completely disregard the fundamental physics.

There is a grain of truth in this, but we must not be misled. First, the physics of our world might be vastly different and yet (for whatever reason) the pattern of heads- and tails-outcomes in HM might be exactly the same; in which case, the chances could be the same. Imagine a universe in which matter is a continuum and obeys something like the laws of Cartesian physics; imagine that gases exist in such a world and spread in the way we are used to. Despite the basic physics being very different, the chance for TD behavior could be

[6] Notice that the replacement of one language by another is not even a requirement for reduction, which only requires that the terms of the two theories be connected by bridge laws. For a discussion of this aspect of reduction see (Dizadji-Bahmani, Frigg & Hartmann, 2010).

the same because the overall pattern of gases spreading in HM might be the same. Second, insisting, as we do, that macro-level chances are genuine chances does not imply that such chances be logically independent of chance rules or laws at the micro-level; i.e., does not imply that such chances cannot potentially conflict with rules and laws at the micro-level. There can be a clash between rules, and as we have seen earlier, we do want macro- and micro-rules to be consistent (we interpreted Equation 3 as a consistency constraint!). But this does not imply that macro-chances are redundant, or that they aren't chances at all (as we saw in chapters 3 and 6, specifically sections 3.4 and 6.2.2).

Residual unease about macro-probability rules and macro-chances might come from the common intuition that there is something epistemic about our probability for a gas to behave TD-like—after all, the gas has one and only one initial condition, and given this initial condition it is determined whether or not it behaves TD-like. This worry can be defused by recalling the discussion of chance and determinism from chapter 3. Information about the precise initial condition (IC) of a particular gas is certainly *inadmissible*, for application of PP, for an agent whose background knowledge includes the micro-level laws: together, such information logically implies how the gas will behave and hence provides (in principle!) knowledge about the gas's behavior that does not come by way of information about the HOC chances. The crucial point is that in typical situations in which we observe gases, we just don't have the inadmissible IC information, or the ability to calculate what 10^{23} particles will do from a given IC, and that is why we use chances and PP to set our degrees of belief. So we use chances when we lack better knowledge. But this does not turn HOCs into Bayesian ignorance probabilities. What it shows is that demons who can observe precise initial conditions and do calculations fast enough have no use for HOCs; they have better information with which to guide their credences about future events than the information HOCs provide. We humans, alas, never have had nor will have either such information about initial conditions, or such demonic calculational abilities. For us, it is a good thing that objective chances exist, and that we can come to know (and use) them.

7.2.3. Summing Up SM

It is plausible that a Best System for our world will capture the probabilities of Boltzmannian SM. In fact, we suggest that it may do so twice over, so to speak: first, in terms of strictly macro-level probabilities for macro-state transitions or for TD-like behavior; and second, by including a chance rule corresponding to the PHPP. Whether the latter truly deserves to be part of the Best System, and is compatible with the former in the sense demanded by HOC, is something that cannot be proven. But we feel that the successes in application of SM so far provide positive evidence for this.

7.3. Quantum Chances

7.3.1. "Standard" QM

As we noted in the introduction to this chapter, it was with the rise of QM that physicists first started to take seriously that the fundamental physics of our world might be irreducibly probabilistic, that is, able to provide at most probabilistic predictions concerning future events, even given a *complete* description of the state of affairs at a certain moment of time. While this turn away from determinism was unwelcome to many physicists, Einstein most prominently, others embraced this and the other radical oddities of the quantum theory that emerged in the 1920s. The radicals (Bohr, Heisenberg, Pauli, Jordan, and others) eventually won out over the skeptics,[7] leading to the establishment of a standard formulation of QM and an accompanying interpretation of the formalism that remain, nevertheless, controversial even today. The cause of continuing controversy is not the fact that QM introduces fundamental probabilities per se, but rather the way in which it does so: in the context of *measurements*.

[7] For discussion of the history of how the "orthodox" interpretation of quantum theory came to predominance in physics, see (Cushing, 1994).

Although we cannot give a complete introduction to the basic features of QM here, we will briefly recount some of the basic elements of standard QM and how they lead to the so-called measurement problem, which is crucial for an understanding of how probability does and doesn't function in standard QM.[8]

QM describes physical systems by specifying a mathematical quantum state, represented by a vector in some vector space. This vector space, often called a Hilbert space, permits the representation of all the possible physical states that the system might have, given the kind of system it is (e.g., electron, or proton, or a pair of electrons emitted from a single source, or an *He* atom, etc.). A curious and important feature of QM is that there is more than one way to represent a system using vectors in a vector space, and the different ways correspond to different *properties* that the physical system can be seen as having: properties such as *total energy*, or *momentum* (in some frame of reference), or *position* (again, in some coordinate frame).

The representation of microscopic systems in terms of *position* and *momentum* is important for many applications of QM. However, properties that can (in principle) take on a continuum of distinct values, such as position (and momentum) involve—even for quantum systems that consist of only a single particle—an *infinite*-dimensional Hilbert space, and bring with them certain mathematical complexities that are not relevant for understanding the core notions of QM and the measurement problem to which they give rise. So we will introduce QM with the example of the spin of a simple quantum particle, such as an electron or proton, which has only a two-dimensional Hilbert space.

Spin is genuine quantum property that has no real classical equivalent. The closest "classical equivalent" one can find is the angular momentum of a sphere rotating around an axis that goes through its center (like the rotation of the earth around its axis), where the magnitude of the angular momentum depends on the radius of the sphere, its mass, and the speed of rotation. But thinking of the quantum spin

[8] Excellent introductory texts to conceptual issues in QM are (Albert, 1992) and (Maudlin, 2019). For an in-depth yet non-expert introduction to the formalism of QM see (Hughes, 1989).

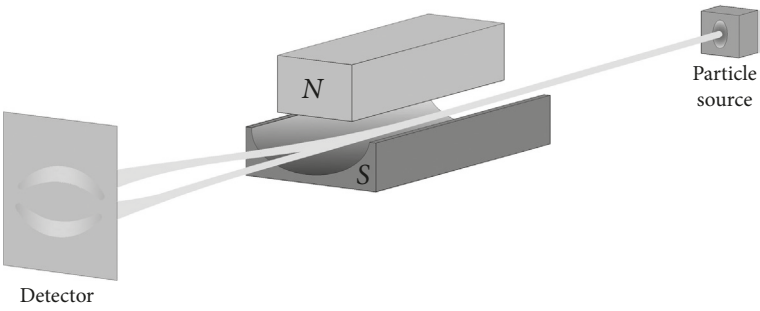

Figure 7.4. Stern-Gerlach device for measuring electron spin.

of an electron in classical terms is misleading because the electron is a point particle (and as such has no radius) and it makes little sense to say that a point rotates.[9] So it is best not to visualize spin and instead think of it in terms of its experimental manifestations. The most important of these manifestations is what we observe when an electron is placed in an external magnetic field. QM associates a magnetic moment with spin and hence when an electron is placed in a magnetic field it experiences a force. The effect of this force is observed in a famous experiment called the Stern-Gerlach experiment. Figure 7.4 shows a schematic representation of the experiment.

The core of the experiment is a magnet that produces an inhomogeneous magnetic field. A source (shown on the right-hand side) produces a beam of electrons that moves through the magnet. While in the magnet, magnetic moment of the spins interact with the magnetic field, either pulling the electrons up (if the north pole of their magnetic moment points downward) or pushing them down (if the north pole of their magnetic moment points upward). As a result, the electrons deviate from the straight path they would have taken in the absence of the magnetic field. A screen at the end of the magnet (shown on the left-hand side) registers where the electrons

[9] From the perspective of later quantum field theories, the electron is not a point particle, or indeed a particle at all, but rather a "quantum of excitation of the quantum field." In so far as we understand what this means, it seems clear that it remains true that it makes no sense to think of such a thing as rotating.

are. Let's assume we have a Cartesian coordinate system that is such that the beam moves along the *x*-axis and the vertical direction is the *z*-axis. Classically one would expect that the magnetic moments in the *z*-direction could have any magnitude, and so one would expect to find a vertical line on the screen. However, the experiment shows two points rather than a line. So all electrons are diverted either upward or downward by the same amount, and so in effect the beam gets split into two. There is never any other diversion, and in particular no zero diversion.[10] Moreover, the beams going up and down have the same intensity, and so we infer that the same number of electrons are diverted upward and downward. From these observations, one infers that the angular momentum of an electron in *z*-direction can have only two values (which have same magnitude but opposite signs), and that these values occur with the same frequency in the beam emitted by the source.

QM explains this experimental behavior of electrons by postulating that an electron can have two *z*-spin states, $|u\rangle$ and $|d\rangle$. The brackets indicate that the states are elements of a vector space, the system's Hilbert space, and the *u* and *d* are mnemonic notations motivated by the motion of the electron in the magnetic field: *u* stands for "up" and *d* for "down." If the electron is in state $|u\rangle$, it moves upward when moving through the magnetic field; if it is in state $|d\rangle$, it moves downward. This might suggest that the Stern-Gerlach experiment could be interpreted as a device that *reveals* the quantum state of an electron: the source produces electrons that are either in state $|u\rangle$ or $|d\rangle$, and the apparatus sorts the original mix of both into two beams, one in which all electrons are in state $|u\rangle$ and one in which all are in $|d\rangle$.

Unfortunately, this is not what happens and things are a bit more complicated—and these complications matter because the moves that QM makes to circumvent the complications is what brings the probabilities into the theory. Quantum states, as mentioned earlier, are elements of a vector space. This matters here because it implies that they can be multiplied with constants and added together, and

[10] This is so for particles with half-integral spin (±1/2, ±3/2, ±5/2, etc.). Electrons and protons are such particles (fermions). Bosons, however, have integral spin and can have spin values of zero, which can lead to no diversion.

the result of such an operation is still a member of the vector space. Specifically, if $|u\rangle$ and $|d\rangle$ are in the vector space, then $|v\rangle = a|u\rangle + b|d\rangle$ is also in the vector space for any complex numbers a and b that satisfy the condition $|a|^2 + |b|^2 = 1$. For this reason, the electron cannot only be either in state $|u\rangle$ or $|d\rangle$; it can just as well be in any state of the form $|v\rangle = a|u\rangle + b|d\rangle$. In QM jargon, such a state is referred to as a "superposition." An electron being in a superposition is more than just a "theoretical," or even "mathematical," possibility—such states do occur in nature. In fact, it turns out that the electrons that are produced by the source are in a state like $|v\rangle$ (and there are experimentally established ways to *prepare* a stream of electrons so that they can be known to be in a state like $|v\rangle$, if one is unsure about their initial state. (For a detailed account of one way to do this, see (Albert, 1992), Ch. 2).

But now we are in trouble. What outcome are we to expect in the Stern-Gerlach experiment if the electron enters the magnet in a state like $|v\rangle$? Our current interpretation does not cover such cases, and there does not seem to be room for states like $|v\rangle$ in the account as developed so far. If the electron is in state $|u\rangle$ it moves up and if it is in $|d\rangle$ it moves down, and these are the only two things that happen in the experiment. So there is no experimental outcome we could assign to $|v\rangle$. This problem is resolved by the infamous *Collapse Postulate*, which in the current context says that when z-spin measurement is performed, then the state of the electron collapses either to state $|u\rangle$ or $|d\rangle$, no matter what the electron's state prior to measurement is. That is, the postulate says that if an electron in state $|v\rangle$ enters into the measurement device, the electron will be either in state $|u\rangle$ or $|d\rangle$ when the measurement is concluded. So the z-spin measurement device changes the electron's state to a state that is interpretable in terms of z-spin properties!

The remaining question is which of the two possibilities is actualized. What standard QM says is that the outcome is not determined in advance, but is rather a matter of chance. This is codified in the *Born Rule*, which assigns probabilities to the different possible outcomes of a measurement. In our context the rule says that if an electron in state $|v\rangle = a|u\rangle + b|d\rangle$ enters the measurement device, then the probability

of finding the electron in either of the two z-spin states is equal to the square of the coefficient of this state in $|v\rangle$. Specifically, the probability of finding an electron in state $|u\rangle$ is $|a|^2$ and the probability of finding it in state $|d\rangle$ is state $|b|^2$. So if the electrons enter the system in state $|s\rangle = \sqrt{1/2}\,|u\rangle + \sqrt{1/2}\,|d\rangle$, then each electron has a probability of 1/2 to come out in state $|u\rangle$ and a probability of 1/2 $|d\rangle$. A beam contains an extremely large number of electrons, and for this reason about half of the electrons end up in state $|u\rangle$ after the measurement, and the other half in state $|d\rangle$. This explains the splitting of the beam into two rays of equal intensity.

It's worth emphasizing that the preceding is merely an example, and the two rules we have applied to the case—the Collapse Postulate and the Born Rule—are completely general. Assume we perform a measurement of property P, and let us call "P-states" the states that can be interpreted as giving definite P-outcome (in the previous example P is z-spin, and $|u\rangle$ and $|d\rangle$ are z-spin-states). Whenever we perform a measurement of property P and the system's state prior to measurement is in a superposition of P-states, then the state collapses onto one of the P-states with a probability equal to the square of the coefficient of this state (in our example the probabilities are $|a|^2$ and $|b|^2$, respectively).

These two rules are immensely effective in generating successful predictions. To date no experiment is known that contradicts QM as formulated in the preceding, and it is the version of QM that working physicists operate with. For this reason we call it *standard* QM. But qualifying a theory as "standard" is meaningful only if there are "non-standard" versions of it. And there are. To understand what these versions are, we need to articulate the (in)famous *measurement problem* of QM, which arises in the preceding theory. While all versions of QM retain (some version of) the Born Rule, the collapse postulate is highly controversial and the main aim of non-standard versions of QM is to exorcise the postulate.

We cannot review alternative versions of QM here, but we would like to explain at least briefly why the Collapse Postulate is widely considered to be problematic. Like CM, QM has dynamical rules (or laws) describing how the state of any system should evolve over time.

In CM states evolve according to Newton's (or Hamilton's) equation of motion, and in QM they evolve according to Schrödinger's equation of motion. QM is a general theory and so its rules for the time evolution of states in principle cover how *any* material system should evolve over time; the system could be a single particle subjected to some potential, or a set of particles, or even a huge collection of particles. A fortiori they also cover what happens when a pair of distinct systems, initially separate from each other, come into interaction with each other. A particular case of such an interaction is when a small quantum system (such as an electron) interacts with a large multi-particle system (such as a Stern-Gerlach magnet); that is, the kind of interaction we tend to think of as a *measurement*. Now, an essential mathematical feature of the dynamical laws of QM is their "linearity." We need not worry about how linearity is defined here, because its consequence is easy to state non-mathematically. When a small quantum system in a *P*-property superposition state interacts with a *P*-measurement device, what the dynamical laws entail is that the superposition should "infect" the device: the combined system should evolve into one in which the small quantum system is still not in any definite *P*-property state, and the device is in a superposition of distinct measurement-result states! In other words, what the dynamical rules alone entail is that measurements do not have single results, but rather *all* the possible results. The quantum description of the combined system becomes a superposition, and there is no such thing as "the result that actually occurs." This does not, however, appear to be what really happens when we do measurements in our labs with things like a Stern-Gerlach apparatus. Reality appears to keep our Stern-Gerlach apparati in one definite state, corresponding to one or the other of the possible *z*-spin measurement results.

One version of the measurement problem, then, is that the dynamical laws of QM predict no definite outcomes, and instead predict that macroscopic objects should evolve into superposition states, but this does not appear to correspond to our experience. But the reader has already seen the solution to this problem that standard QM offers: the Collapse Postulate! To use the theory in the context of understanding measurement interactions, one sets aside the standard laws of dynamical evolution at the point where one considers a measurement to have

occurred, scratches out the superposition, and rewrites the quantum state of the system as the definite-*P*-property state corresponding to the observed measurement outcome. Similarly, the quantum description of the measuring device (if one were interested in thinking about it, which usually one is not) should be not a superposition, but a nice definite-*P*-property-registering state. The collapse rule saves QM from giving us an apparently wrong description of what happens in our laboratories, and allows the theory to make exceptionally well-verified predictions regarding measurements of all sorts.

The price of this solution of the measurement problem, however, is high—for many, too high to be borne. This solution involves introducing an *epistemic* notion, *measurement*, into the dynamical laws of a fundamental theory in an essential way. From a physical standpoint, a measurement device is just a many-particle system, no different from any other system with a similar number of atoms. A puddle or a vapor cloud can deviate the path of a flying electron, just as a Stern-Gerlach apparatus does; what makes the latter interaction dynamically special? For that matter, what *is* a measurement? Which interactions count as measurements and which do not—and *why*? Although it is clear enough how to apply QM in all practical situations, the need to introduce the notion of measurement into the fundaments of the theory leaves it ill-defined in a way that most other physical theories are not, leading many philosophers and physicists to look for some revised interpretation of the theory, or even a replacement for it, that removes the need to talk about measurements as a special kind of interaction. To date, no alternative to the standard theory has gained anything close to acceptance by a majority of physicists (or philosophers). For this reason, and to keep the discussion manageable, we discuss probabilities as they occur in standard QM.

7.3.2. Standard QM and HOC

Thus, in the 1920s a remarkable situation came into being, one that persists to this day: our best fundamental theory of ordinary matter does not give us a coherent, realistically understandable *description*

of what matter is and how it behaves, but rather only an instrumental *prescription* for making predictions—probabilistic predictions, in the main—about the results of observations and experiments. As Tim Maudlin so eloquently argues in his (2019) introduction to QM, one should really not think of QM as a *theory*, but rather as a *recipe* for making certain sorts of predictions.

And as everybody knows, this instrument, this recipe, is fantastically successful: the patterns of measurement outcomes in actual events that constitute quantum experiments are just as you would expect given the objective probabilities generated by the Born Rule. Another way to put this fact: the Humean Mosaic of our world appears to have certain widespread, reliable stochastic-regularity patterns in it, patterns that can be captured and systematized with great strength, amazing simplicity, and fantastically good fit, by the axioms of standard QM including the Born Rule. In other words: prima facie, standard QM is a strong candidate for inclusion in the Humean Best System for our world, *even* (nay—especially!) *if taken as nothing more than a recipe for predicting events that are all, in the end, human observation-events*.[11]

This is so, however, only for a pragmatic HOC such as we advocate in this book, and apparently not for Lewis' Best System theory, because in the latter the laws are supposed to only invoke the *perfectly natural properties* (natural kinds) from which events in the fundamental HM are composed. Whatever account of "perfectly natural properties" one may offer, the property of *being a measurement* is presumably not going to count as perfectly natural. By contrast, a pragmatic approach such as ours, which is happy to look for patterns in the HM at various "levels," expressible with a wide range of predicates and terms that have proven useful for human beings—trains and tables as

[11] By judiciously ignoring the quantum measurement problem, the apparent limitation of this recipe to predicting observations made by humans and other rational beings can be pushed out of sight, and one can then talk indiscriminately about events such as radioactive atoms decaying, electrons detaching from an atom by quantum tunneling, and so forth, independent (apparently) of whether measured/observed or not. Physicists do this all the time, and it is "cheating," for reasons well understood in the philosophy of QM literature. But even if we avoid cheating and stick to merely systematizing actual human measurements and observations that are covered by the QM recipe, there is still an awful lot of stuff in the HM that gets simply and strongly systematized by a Best System that includes the QM rules.

well as quarks and electrons—need not blush at incorporating a rule such as the Born Rule, even if it refers to a kind of event that can only be understood as the activity of an epistemic agent.

So we have eminently good reason to think that a Best System of our world, under the pragmatic account of HOC defended in this book, would include the Born Rule as an element, with the rest of the theory of standard QM written into the system as auxiliary content necessary to specify the chance setups in which the rule applies and to determine the mathematical values of the probabilities. There is, moreover, no reason why this should be restricted to the non-relativistic QM that emerged in the late 1920s. Quantum field theories, which now are systematized in the Standard Model of "particle" physics,[12] share all the relevant features of QM just discussed, and would deserve to be parts of the Best System for exactly the same reasons.

Can alternative views of objective chance also make good sense of QM probabilities? As we just noted, on the face of it, Lewis' account is difficult to square with the Born Rule's invocation of *measurement*.[13] For similar reasons, chance primitivists ought to feel uncomfortable with standard QM as well: why is it *measurement situations* that have associated objective chances for outcomes, rather than, say, *interactions of big/many-particle systems with little/few-particle systems*? The latter is awfully vague, but at least it is not overtly epistemic. And many physicists and philosophers of physics have speculated that something like this must really be the better description of where nature goes indeterministic.

The idea would be that the big-system/little-system interaction somehow provokes an objective "collapse" of the quantum state of the little system, taking it from superposition to eigenstate (or at

[12] We put "particle" in scare quotes here because the majority view among physicists is that *fields* are fundamental rather than particles. We take no sides in this still-contentious debate.

[13] Barry Loewer's version of the Best System theory of laws, however, is explicitly more pragmatic in its approach, and might be able to lay claim to standard QM as a likely part of the Best System for our world. I suspect, however, that Loewer would prefer to see some replacement version of QM (e.g., some relativistic successor to Bohm theory or GRW) built into the Best System. Although more pragmatic than Lewis, Loewer is still aiming to offer an account of the true fundamental physical laws.

least "sharply localized" states, for continuous variables such as position). There are problems with this idea. First, it is no good if left in the vague form indicated; to be a worthy proposal, it would have to be cashed out in some concrete and precise fashion. This would make the resulting theory, in principle, at least *slightly* different from standard QM in its predictions, since standard QM involves no spontaneous collapses and instead says that the superposition of the little system can turn into a superposition of the combined big-system/little-system system (this being the lesson of Schrödinger's cat). It would be a new theory, in other words, and not standard QM. Second, experiments that have gone looking for spontaneous collapses have so far not found any. Nature so far appears stubbornly fond of the instrumentalist QM recipe (though it is clear that, depending on how the spontaneous collapses work, it may be near-impossible to design an experiment that would reveal them). This is the case for the GRW theory, which is precisely a well-worked-out spontaneous collapse theory of the type we are considering here; for a discussion of this theory and how HOC may work for it, see (Frigg and Hoefer 2007).

So standard QM is not a very apt home for primitive chances or chancy fundamental laws, but some nearby, reformed theory (such as GRW) might look like a natural candidate for a primitive-chance reading. That is, it might look that way at first blush; but the arguments of chapter 1 aimed to convince the reader that this interpretation is not nearly as clear and unproblematic as it may initially seem. Readers who feel attracted to such a reading of QM probabilities are invited to re-read section 1.3.

7.3.3. Humean Chance and the Wave Function

We end with a final note about how HOC squares with standard QM more easily than a strict Lewisian Best System approach does. In our introduction to the main features of QM in the preceding, we chose to look at spin properties and states; these are easy to represent and to use to illustrate the notions of superposition and measurement. We did not look at how QM represents either the *position* of particles and

or their *momenta*. Position and momentum, unlike spin, are variables that range over a continuum of possible values, and quantum states of position or momentum are represented by continuous functions—so-called *wave functions*—defined over the continuous spaces that codify all the possible particle positions (or momenta—but from now on, we will set aside momentum and just think about position) of the quantum system.

To be more precise: the wave function ψ of a system is a complex-valued field defined in the so-called *configuration space* of the system.[14] The notion of "configuration space" is borrowed from CM, and is very similar to the notion of phase space we saw in section 7.2. In CM, if we have a system of N particles, it is sometimes convenient to represent the positions of all the particles through a unique point $Q \equiv (\boldsymbol{Q}_1, \boldsymbol{Q}_2, \ldots, \boldsymbol{Q}_N) \in \mathbb{R}^{3N}$, where $\boldsymbol{Q}_i \in \mathbb{R}^3$ are the position coordinates in physical three-dimensional space of the i^{th} particle. Configuration space is the set of all points that—like Q—represent a possible configuration of all the particles of the system in physical three-dimensional space, and it trivially follows that configuration space is $3N$ dimensional.[15]

In QM, the fact that the wave function ψ of a system is defined in configuration space implies the following. If we have a one-particle system, its wave function assigns a number to each point of the ordinary, three-dimensional space. In this respect, the wave function of a one-particle system can be thought of as analogous to a classical field: at each point in space, the wave function has a certain numerical value, or *amplitude*. If we have a two-particle system, however, then six spatial coordinates—not just three—must be used in order to specify the wave function. And in general, for an N-particle system, $3N$ spatial coordinates are required to specify the wave function.

[14] In other words, the wave function assigns a complex number to each point of that space. The fact that the values of the wave function at points (also called the *amplitude* at that point) is a complex number is a detail that we will ignore from now on.

[15] A difference between configuration space and phase space is that the phase space for a system of N particles has $6N$ dimensions, because a point in phase space represents not just the instantaneous spatial configuration of all the particles, but also their velocities or momenta.

A *point* in configuration space represents each of the *N* particles having some *specific* location in ordinary 3-space. (And a continuous line in configuration space could by extension represent how *N* particles all change their locations in space, over time.) But the wave function does not select a point of configuration space; it assigns some field value, or amplitude, to every point in configuration space. How does such a thing *represent* the state of a system of *N* particles?

This is where probability (again) enters into the story. In this context, the Born Rule says that the wave function's amplitude at a point (or better: the integral of the amplitude over some region) of configuration space determines the probability of finding the positions of the *N* particles to be such as to fall within that region, if we were to measure the positions of the particles. A wave function that is sharply peaked near a certain point of configuration space—that is, with amplitude near-zero everywhere except in a small region centered on that point—represents that if we measure the positions of the particles we will find them to have those specific values ($\pm\varepsilon$), with probability near 1. But a wave function that is smeared out broadly, or has many distinct lumps or bumps of non-trivial amplitude in different regions, represents our set of *N* particles as not having *any* definite positions, at the moment, and merely having certain probabilities of *being found* in certain places, if we go looking for them. The probabilities are, again, given by the Born Rule, which in the current context (roughly) says that the probability of finding the particles in certain locations is the square of the amplitude of the wave function for these locations. In other words, such a wave function represents the particles as being in a superposition state, a superposition of multiple possible precise location-states. Figure 7.5 illustrates this idea for the simple case of a two-particle system, with only one spatial dimension represented for each particle.

Now that we have seen how QM treats position, we are ready to raise our final concern about the Lewis-style approach to Humean chance in QM.

Regarding the wave function, we might wonder how such a thing, if taken as a physically real element of our world, fits into the HM. If we stick to Lewis' line concerning the nature of the HM, it seems

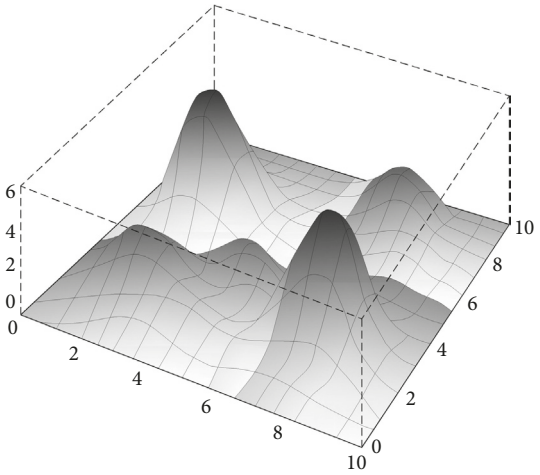

Figure 7.5. Wave function of two particles in one spatial dimension, in an infinite potential well (i.e., a box).

that there is a problem: the wave function cannot be defined in 3-d space, and straightforward attempts to boil its content down into something that *does* fit into 3-d space inevitably lose vital information encoded in the full wave function.[16] But redefining the HM to fit QM is also very problematic. Should it be configuration space that is the basic backdrop for the HM? If so, should it be the configuration space of position, or momentum? If instead we choose an infinite dimensional Hilbert space as the backdrop for the HM, then the universe's quantum state becomes a mere vector or ray singled out in that space;

[16] The fact that one of standard QM's key elements, the wave function, is defined in a high-dimensional space of some sort is a problem for more than just the Lewisian-Humean approach to chance; it is an issue that all those who would like to be realists about QM must confront. One approach to the concern is to try to use the fact (if it is one) that QM is only an effective, non-relativistic-limit approximation to some deeper, relativistic quantum field theory (which is set in ordinary spacetime) as a way to bring the wave function back down to reality. Ways of recovering the wave function from quantum field theory are explored in (Wallace & Timpson, 2010) and (Myrvold, 2015). Given the limitations and interpretive difficulties of quantum field theories themselves, we feel that this way to resolve the ontological issue is promising but far from being clearly workable.

what kind of a mosaic is that, and how can we discern patterns in it that correspond to familiar scientific regularities?

We don't here want to claim that no successful responses to these questions can be found, but we do want to highlight again the advantage that our pragmatic HOC has when it comes to capturing QM chance rules, and particularly those of standard QM. Our inclusive, multi-level understanding of the HM makes it easy to locate precisely the pattern of successful quantum experimental results that did, in actual history, lead to the acceptance of QM as a pattern ripe for systematization, by standard QM understood as an instrumental recipe for making probabilistic predictions. No matter what the real fundamental level of our universe is like, the meso- and macroscopic patterns of events in 4-d spacetime that we are familiar with must turn out to be *real*—to supervene in the right way on that underlying fundamental level—and, hence, to be grist for the pragmatic Best System mill. HOC is, in short, the best theory (among all extant reductive theories of chance) for capturing the probabilities of quantum theory in an elegant fashion.

7.4. Summing Up

Probabilities are found in physics in both deterministic and indeterministic settings. Either variety can be accommodated by HOC in natural and straightforward ways. In this chapter we looked at two of the most paradigmatic physical theories in which objective probabilities play a central role: classical statistical mechanics, and standard non-relativistic quantum mechanics. In the former, we found that the Best System for a world in which SM is a useful theory could incorporate Humean chances in two ways. First, the Humean Mosaic might be such that the pattern of transitions from one thermodynamic macro-state to another give rise directly to Humean chances for such macro-state transitions. Second, more directly, the pattern of actual evolutions of systems well-described by SM—that is, the patterns of how the states of such systems, represented by points in phase space, begin and evolve over time through various macro-states and eventually to equilibrium—could be such as to have a simple and elegant

systematization in the combination of classical mechanics + the Past Hypothesis Proportionality Postulate. In either or both of these ways (if sufficiently compatible), the chances of SM can plausibly turn out to be proper Humean chances.

When it comes to quantum probabilities, we found that the concordance between HOC and the theory is practically a marriage made in heaven: what the theory gives us is, on the face of it, a recipe for calculating probabilities for experimental observations. Since that recipe is enormously well confirmed by actual experiments, we have overwhelming reason to think that the Best System of Humean chances for our world will include that recipe in its compendium of rules. If we ignore the measurement problem and move beyond experiments out "into the wild," and think of QM as prescribing probabilities for such things as decays of radioactive atoms, scattering cross sections for particles moving through various media (e.g., x-rays through the atmosphere, neutrinos passing through the Earth, electrons tunneling past a potential barrier, etc.), we again have overwhelming reason to think that observable events in the HM for our world concord with QM, and so the case for QM being part of the Best System becomes even more compelling. This is a fact that will not change, no matter whether QM is replaced by an ontologically clearer successor theory at some future date or not, and no matter whether that theory is deterministic or not. HOC is in this sense ideally suited to be the right philosophical account of the probabilities given to us by quantum theories.