Typicality versus Humean probabilities as the Foundation of Statistical Mechanics

Dustin Lazarovici*
Université de Lausanne, Faculté des Lettres,
Section de Philosophie, 1015 Lausanne, Switzerland

August 24, 2019

The paper discusses two contemporary views about the foundation of statistical mechanics and deterministic probabilities in physics: one that regards a measure on the initial macro-region of the universe as a probability measure that is part of the Humean best system of laws (Mentaculus) and another that relates it to the concept of typicality. The first view is tied to Lewis' Principal Principle, the second to a version of Cournot's principle. We will defend the typicality view and address open questions about typicality and the status of typicality measures.

1 From microscopic laws to macroscopic regularities

Consider the following macroscopic regularities that we observe in our universe:

1) Apples do not spontaneously jump up from the ground onto the tree.

2) Rocks thrown on earth fly along (roughly) parabolic trajectories.

3) The relative frequency of *heads* in a long series of fair coin tosses comes out (approximately) 1/2.

These regularities are all of a different kind. 3) is a statistical pattern. 2) is a mechanical phenomenon. 1) turns out to be an instance of the second law of thermodynamics. All

*Dustin.Lazarovici@unil.ch

three regularities strike us a law-like; arguably, they are even among the more basic experiences founding our belief in a lawful cosmos. However, it turns out that none of them is nomologically necessary under the fundamental laws that we take to hold in our universe. In fact, given the huge number of microscopic constituents of macroscopic objects[1] and the chaotic nature of the microscopic dynamics (small variations in the initial conditions can lead to significant differences in their evolution), the fundamental laws put very little constraints on what is physically possible on macroscopic scales.

It is possible that particles in the ground move in such a coordinated way as to push an apple up in the air (we know that because the time-reversed process is common and the microscopic laws are time-reversal invariant). It is possible for a balanced coin to land on *heads* every single time it is tossed. And it is possible, as Albert (2015, p. 1) so vividly points out, that a flying rock is "suddenly ejecting one of its trillions of elementary particulate constituents at enormous speed and careening off in an altogether different direction, or (for that matter) spontaneously disassembling itself into statuettes of the British royal family, or (come to think of it) reciting the Gettysburg Address."

Assuming deterministic laws, a physical event or phenomenon is (nomologically) possible if and only if there exist micro-conditions of the universe that evolve under the microscopic dynamics in such a way that the event or phenomenon obtains. Given our limited epistemic access to the micro-state of the universe (or any complex system, for that matter) we thus need some inferential procedure from the fundamental dynamics to the salient macroscopic regularities, other than finding the exact solution trajectory that describes our universe. In fact, even if we *did* know the exact initial conditions and could predict the entire history of the universe deterministically, it would seem odd if law-like regularities as the ones stated above turned out to be merely accidental, contingent on the very particular microscopic configuration of our universe. That is, even if we were Laplacian demons and could verify that dynamical laws + initial conditions make (let's say) the second law of thermodynamics true in our world, we should care for some additional fact or principle that makes it counterfactually robust and gives it more nomological authority.

Some people find it preposterous to refer to initial conditions of the universe in order to account for something like the motion of a rock, or the cooling of a cup of coffee. Well, in practice, we don't. In principle, however, even the best-isolated subsystem is part of a larger system with which it has interacted at some point in the past. Hence (to adopt an expression from John Bell), if we make postulates or assumptions about initial conditions of various subsystems individually, we commit redundancy and risk

[1]The relevant order of magnitude is given by Avogadro's constant which is $\sim 10^{24}$ per mol.

inconsistency. Any attempt at a conclusive and fundamental account must therefore talk about the universe as a whole (cf. Oldofredi et al. (2016)).

In the context of classical mechanics – that we shall focus on for now – there is a reasonably widespread agreement that the following holds true as a mathematical fact:

> There exists a small (low-entropy) region M_{PH} in the phase space Γ of the universe such that the uniform Lebesgue or, more precisely, Liouville measure λ on M_{PH} assigns high weight to initial conditions leading to micro-trajectories that instantiate the thermodynamic regularities – in particular, the *thermodynamic arrow of time* – and other salient patterns (about coin tosses, stone throws, etc.) that we observe in our universe.
>
> That is, if we denote this set of "good" initial conditions by $M^*_{PH} \subset M_{PH}$, it holds true that $\lambda(M^*_{PH})/\lambda(M_{PH}) \approx 1$.

Following recent lectures of David Albert, we shall call this the *fundamental theorem of statistical mechanics* (FTSM). "Statistical mechanics" here can be understood very broadly, as being tasked with explaining or predicting macroscopic regularities on the basis of the microscopic laws. An important question is, of course, why the above statement seems so compelling given that it is virtually impossible to prove in any rigorous sense. This, however, is not the focus of the present paper, so suffice it to say that (if nature is kind to us) arguments going back to Ludwig Boltzmann strongly support the FTSM or a suitably close variant (see e.g. Bricmont (1995); Penrose (1999); Albert (2000); Carroll (2010); Goldstein (2012); Lazarovici and Reichert (2015) for detailed discussions). What our present discussion is going to focus on instead, is the physical and philosophical interpretation of this result – assuming it holds true as a mathematical fact –, as well as the meaning and status of the measure figuring in it.

David Albert (2000, 2015) and Barry Loewer (2007, 2012b) have developed a popular and well-worked out position in the context of the Humean Best System Account of laws (BSA), adapting David Lewis' theory of objective chance (Lewis, 1980, 1994; Loewer, 2001, 2004). According to their proposal, the best system laws of our world consist in

1. The deterministic microscopic dynamics.
2. The Past Hypothesis postulating a low-entropy initial macrostate of the universe.
3. A probability measure $\mathbb{P} = \frac{\lambda}{\lambda(M_{PH})}$ on the Past Hypothesis macro-region M_{PH}.

This probability measure does not refer to any intrinsic probabilities or random events in the Humean mosaic. Its inclusion into the best systematization is justified by the fact that it comes at relatively little cost in simplicity but makes the system much more informative, precisely because it accounts – via the FTSM – for the thermodynamic 'laws', the entropic arrow of time, and many other macroscopic regularities.

Loewer introduced the name "Mentaculus" for this best system candidate, a self-ironic reference to the movie *A Serious Man* in which an evidently crazy person tries to develop a "probability map of the entire universe". As a philosophical proposal, though, the Mentaculus is not completely crazy, as it attempts to provide a precise account of deterministic probabilities which avoids the problems of standard subjectivist and frequentist interpretations. Moreover, Albert (2000, 2015) and Loewer (2007, 2012a) employ the Mentaculus in a sophisticated analysis of counterfactuals, records, and special science laws, the details of which are beyond the scope of this paper.

Another view that has been defended by some authors (e.g., Goldstein (2012)) is that the Liouville measure on the initial macro-region should be understood as a *typicality measure*. This is to say that the FTSM is interpreted not as a probabilistic statement but as the proposition that the macroscopic regularities in question obtain in *nearly all* or *the vast majority of* possible worlds (consistent with the dynamical laws and the Past Hypothesis). In this sense, macroscopic regularities – such as the second law of thermodynamics – come out as *typical regularities* (Lazarovici and Reichert, 2015).

In general, a property P is *typical* among a reference class W if nearly all members of W instantiate P. The basic claim is then that, when applied to a reference class of nomologically possible worlds, typicality facts can serve an explanatory, epistemic, and behavior-guiding function – maybe more convincingly so than probability facts (Goldstein (2001); Maudlin (2007); Volchan (2007), Hubert (forthcoming)).

I will not provide a comprehensive exposition of typicality right away. Instead, we will develop the position out of a critique of the Mentaculus. The goal is to argue that it is preferable to adopt typicality even in the context of a best system account, while there are additional motivations for it if one holds an anti-Humean view about laws of nature.

2 Principal principle and the meaning of Humean probabilities

Typicality facts are distinct from probability facts (Wilhelm, 2019). The first and most obvious difference between probability and typicality is that the latter is not a quantitative notion. Typicality statements may be context-dependent (as we will discuss in section 7), but in any given context, a property of the world can only be *typical* or *atypical* (or neither), though not "more typical" or "less typical" than another. The role of a typicalility measure is only to determine "very large" and "very small" sets (of initial conditions, i.e. possible worlds), giving precise meaning to the notion of "nearly all". The Humean probability measure, in contrast, is supposed to contain much more information. In fact, it will assign a probability – or conditional probability – to any physical proposition about the world (through the set of microconditions for which the event obtains, cf. Albert (2015, p. 8)): a probability that my dog gets sick if he eats a piece of chocolate, or that your favorite football team wins the next Super Bowl, or that the United States elect a female president in 2028.

The epistemic and behavior-guiding function of these predictions is supposed to be manifested in a normative principle, the *Principal Principle* (PP), which states that we should align our initial credences with the objective Humean probabilities. Formally:

$$C(A \mid \mathbb{P}(A) = x) = x, \tag{1}$$

or, for conditional probabilities,

$$C(A \mid B \wedge \mathbb{P}(A \mid B) = x) = x. \tag{2}$$

There are other variants of the PP proposed in the literature, and debates about what constitutes "admissible information" that one can conditionalize on (Hall (1994, 2004); Lewis (1994); Loewer (2004)), but these subtleties will not be essential to our discussion.

In any case, stipulating the PP does not explain *why* it is rational to follow it and what physical information a probabilistic prediction of the Humean best system contains. What exactly is the Mentaculus telling us by assigning, let's say, a (conditional) probability of 30% to the United States electing a female president in 2028? After all, the Lewis-Loewer theory agrees that there are no genuinely probabilistic facts in the world. Every possible event either occurs or not, and whether it does is entailed by initial conditions and the deterministic dynamics. So what do such single-case probabilities even refer to?

A standard Humean response is that, by definition, the probability measure figuring in the best system laws is the optimal measure for our world in terms of balancing simplicity and informativeness. Hence, while a single-case probability may not express anything about the individual event *per se*, there is something about the structure of the Humean mosaic as a whole that makes the particular value true or accurate (Lewis, 1980). Indeed, according to the BSA, $\mathbb{P}(A) = x$ is true in all and only those worlds whose best system implies $\mathbb{P}(A) = x$, so the proposition seems to be saying *something.*

I submit that we cannot get out more of the best system than we put in. The Humean probability law can only inform us about the features or regularities of the world that it is supposed to fit in the first place. If it is more accurate to assign a chance of 30% than of 60% (let's say), there must be concrete physical facts in the world that make it so; and these facts must be among those that go into evaluating the strength of the best system candidates. However, as I will argue in more detail below, many probability measures would assign a probability close to 1 to the thermodynamic regularities and other salient statistical patterns, yet a number very different from 0.3 to the United States electing a female president in 2028. By some standard, these measures may not be as simple as the Liouville measure – which is why they are not part of the best system –, but unless they fare worse in terms of *fit*, there is nothing in the world that makes them less accurate when it comes to predicting presidential elections.

Some authors have read Lewis as suggesting that the Humean probability law is supposed to fit the macro-history of the world by assigning as high a probability as possible to any event that, in fact, occurs, and as low a probability as possible to any event that, in fact, does not occur (while being constraint by the requirement of simplicity). It cannot really work that way. Being a good predictor of presidential elections does not gain you as many points for "fit" as predicting the increase of entropy in our universe. Also, assigning a probability of $1/2$ to individual coin tosses – which may look like the laws are completely ignorant about the outcomes – is actually informative because it implies a probability close to 1 for the event that the relative frequency of *heads* and *tails* in a long series of tosses is approximately $1/2$. In the end, the best system probability law will be one that informs us about robust regularities and global patterns in the world, while the fit to singular events will count little to nothing in the trade-off with simplicity.

In particular: if, given the dynamical laws and the Past Hypothesis, two probability measures $\mathbb{P}$ and $\widetilde{\mathbb{P}}$ are necessarily equivalent in terms of fit – because they predict the same global patterns – while $\widetilde{\mathbb{P}}$ loses out in terms of simplicity, there is *no possible world* in which $\widetilde{\mathbb{P}}$ replaces $\mathbb{P}$ as part of the best system. (That is, unless the metric of simplicity is oddly contingent in a way that depends on microscopic details or isolated macro-events.) Therefore, a proposition like "according to the Mentaculus, the probability of event A is $\mathbb{P}(A)$ (rather than $\widetilde{\mathbb{P}}(A)$)" cannot restrict the set of possible worlds any further than to those instantiating the regularities on which $\mathbb{P}$ and $\tilde{\mathbb{P}}$ actually agree.

To summarize, in other words: According to the Humean view, there are certain "chancemaking patterns" Lewis (1994) on which a probability law supervenes, while "probabilities" for a great many other events – and, in fact, for all measureable sets of micro-states, most of which do not correspond to any meaningful macro-event – come out as a by-product. This by-product, to play on a metaphor by Albert (2015, p. 23), is not a gift from God ("I give you my most efficient summary of the regularities, and you get rational credences for all conceivable events for free") but mostly mathematical surplus; the numbers could be very different, yet the physical content of the law the same. Sure, we can *postulate* that we should live our lifes according to whatever numbers the Mentaculus spits out, but as so often with articles of faith, there is no rational prospect of reward in this world.

2.1 Law of large numbers and typical regularities

In a nutshell, Humean probabilities are supposed to be efficient summaries of regularities. In the Mentaculus account, they turn out to refer, first and foremost, to a measure on sets of possible micro-conditions. What has one to do with the other? In all but few especially nice cases: nothing at all (is exactly my point). In the few especially nice cases, the connection between a statistical pattern instantiated by a series $S = (A_i)_{1 \leq i \leq N}$ of similar events (e.g. a long series of coin tosses) and the probability $\mathbb{P}(A_i) = p$ of the individual events that make up the pattern (e.g. the i-th toss resulting in *heads*) is provided by a *law of large numbers* (LLN), that is, a result of the form

$$\mathbb{P}\left(x \in M_{PH} : \left| \frac{1}{N} \sum_{i=1}^{N} \chi_i(x) - p \right| > \epsilon \right) \propto \frac{1}{\epsilon^2 N} \approx 0. \tag{3}$$

Here, $\chi_i(x)$ is the indicator function mapping each possible initial micro-condition x to 1 if the event A_i occurs, and to 0 if the event A_i does not occur for the micro-trajectory with initial condition x (see Fig. 1).

We can read equation (3) as: "the measure of the set of initial conditions for which the relative frequency of occuring events deviates significantly from p is very close to 0". As a mathematical theorem, stating sufficient conditions for (3), the law of large number requires that the events are in some sense independent or uncorrelated, which is often intuitively compelling but nearly impossible to prove (and may fail much more often than we think). The standard derivation would also make use of the fact that p comes out as the theoretical expectation of the "empirical" distribution $\frac{1}{N} \sum\limits_{i=1}^{N} \chi_i(x)$ with respect to $\mathbb{P}$. In the end, however, the role of the measure in (3) is merely to tell us that a particular set of initial conditions – the initial conditions that lead to significant deviations from the statistical pattern of interest – has a measure close to zero. And at this point, it really doesn't matter where the number p came from and whether we thought about it as a "probability" in the first place. In fact, if (3) holds for the measure $\mathbb{P}$, an analogous statement will hold true for many other measures that agree on the smallness of that set (and, in general, of various other sets related to other statistical regularities), even if they disagree on the values assigned to the events A_i individually. And it will hold true in many cases in which the standard assumptions for the LNN are not satisfied. (For most physical applications, these are much too strong anyways, which is why statistical mechanics is so darn hard.)

Philosophically, it is thus unnecessary and misleading to think of (3) as a *consequence* of the single-case probabilities determined by $\mathbb{P}$. It is really the other way round: What a law of large numbers result does, in effect, it to reduce theoretical probabilities to *typical frequencies.* These typical frequencies, I claim, are all that the fundamental laws can or need to inform us about. In particular, if our best theory tells us that (with near certainty) roughly 1/2 of the coin tosses result in *heads*, I can see how, on this basis, we can begin to justify the rationality of assigning credences about individual tosses accordingly; For instance, by appealing to dutch-book arguments (if I accept bets of less than 2:1 on each one of these event, I can be almost certain to lose money on the long run) or maybe by invoking a principle of indifference with regard to the individual event in the pattern that we are about to observe (Schwarz, 2014).

Ultimately, this view is not tied in any way to Humean metaphysics; though when combined with the regularity theory, it puts the latter back on its feet, by actually reducing probabilities to patterns in the mosaic that the best system could reasonably predict.

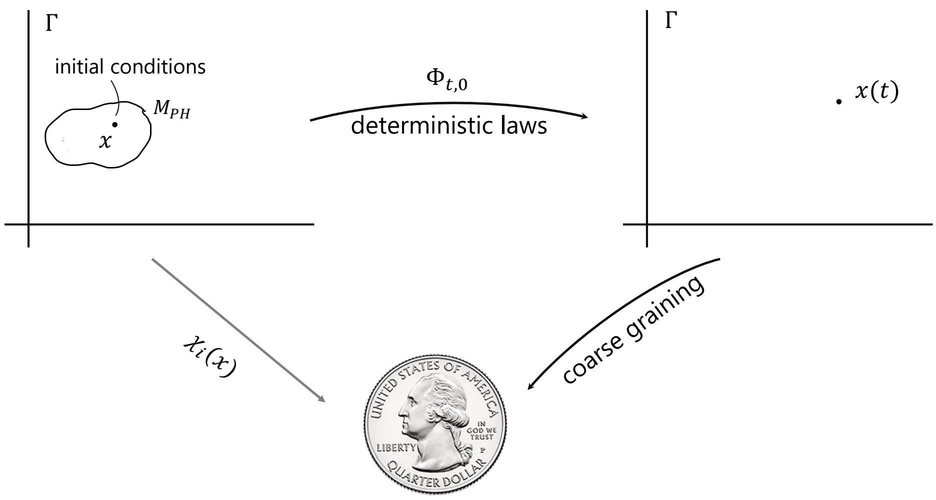

Figure 1: Sketch: a macroscopic event supervening on the microscopic evolution. $\Phi_{t,0}$ is the flow arising as the general solution of the microdynamics.

What about the probability of, let's say, event A: My dog gets sick, given B: He eats the piece of chocolate that I dropped on the floor?

It seems natural to embed this sequence of events into a statistical ensemble: this and that fraction of dogs (of a certain size) get sick if they eat this and that amount of chocolate. If nature is kind to us, the best system will predict this pattern as a typical frequency – which may *or may not* coincide roughly with the conditional Humean probability $\mathbb{P}(A|B)$. It also seems possible to decompose the events into a more fine-grained description which is part of a statistical pattern, e.g., the rate at which a dog's intestinal tract can metabolize theobromine or, finer still, interaction rates of certain molecules.

I believe that the intuition that singular macro-events could or should have an objective physical probability, in addition to a deterministic micro-description, comes from such (in principle) possibilities of embedding or decomposition – which are, notably, non-unique (cf. Hájek (2006) on the reference class problem) and always require further context and analysis.It may be more difficult to explain, in physical terms, what one could mean by "the probability that the United States elect a female president in 2028", but as a starting point, we should look at the statistical regularities (e.g., the sampling methods used to obtain polling data) that such predictions are actually based upon. The idea that the Mentaculus provides a shortcut from the fundamental laws of physics to specific chance prescriptions for individual event token may be philosophically appealing but is ultimately too simplistic.

In the end, whenever we succeed in grounding rational credences in the fundamental laws of physics, they will be grounded in typicality facts. We shall say more about how such typicality facts can be understood. For now, a remarkable point worth emphasizing is that rational credences are thus grounded not in propositions about which our best theories are unsure or uncommital, but in patterns and regularities that they predict, or at least explain, beyond reasonable doubt.

3 Principal Principle versus Cournot's principle

What the previous discussion has, in fact, accomplished is to reduce the Principal Principle (PP) – at least those instances of the PP that could have a basis in physics – to a (weak) version of Cournot's principle (CP). Cournot's principle currently experiences a bit of a revival in philosophy of science (see e.g. Hubert, forthcoming) but has a long tradition in the philosophy of probability, with some version of it being endorsed by Kolmogorov, Hadamard, Fréchet, Borel, among others (see Martin (1996); Shafer and Vovk (2006) for excellent historical discussions).

One way to introduce CP is as a remedy to the following dilemma:

1. Only probabilistic facts follow from probability theory.
2. There are no genuinely probabilistic facts in the world. Any possible event either occurs or not.
3. No facts about the world follow from probability theory.

Assumption 2 could be denied by admitting something like propensities into the physical ontology, but then replace "(physical) facts" by "empirical facts," and we end up with a very similar problem: logically, no empirical facts follow from probabilistic ones.

The Cournot principle can thus be understood as a sort of "bridge principle" leading from probabilistic results to physical/empirical predictions. An unfortunate historical fact is that the formulation provided by its namesake sounds plainly wrong, at least to modern philosophers of science:

> "A physically impossible event is one whose probability is infinitely small. This remark alone gives substance – an objective and phenomenological value – to the mathematical theory of probability." (Cournot, 1843)

Arguably, Cournot understood "physically impossible" in more of an FAPP[2] sense, but the language here has not helped the acceptance of his philosophy.

A more appropriate formulation of the principle can be found in the work of Kolmogorov (1933, Sec. 2.1) : if an event has very low probability, *"then one can be practically certain that the event will not occur"*. Roughly equivalent (by controposition): if an event has very high probability, one should expect it to occur. Other authors have cast the principle in more decision-theoretic terms; in a nutshell: it is rational to act as if very high probability outcome(s) will obtain.

Tim Maudlin (private communication) has convinced me that the most pertinant formulation, in the physical context, is in terms of *explanation* (see also Lazarovici and Reichert (2015), Wilhelm (2019), Hubert (forthcoming)). If a feature of our world turns out to be typical, according to a theory, we should consider it to be explained by that theory. If we observe a robust phenomenon that turns out to be atypical according to our theory, we should look for further explanation – and possibly, in the last resort, revise the theory if none can be given.

In the end, I believe, these are all but different aspects of the same rationality principle which may be cast in terms of expectation, belief, acceptance, explanation, etc. Working

[2]For all practical purposes

out the subtle differences and interrelations between these concepts is certainly worthwile but beyond the scope of the present paper. However, it is important to keep in mind that when we speak about "belief" or "expectation" in the present context, we mean believes or expectations grounded in inferences from our fundamental theories of nature. They thus come with certain normative implications for which phenoma require (further) explanation and which theories can be accepted.

We should also note that while some version of CP could be regarded as a special case of PP, it is characteristic of the view associated with Cournot to privilege statements of "very high" or "very low" probability with respect to their physical or empirical content. Thus Borel's famous (though still somewhat too strongly worded) credo: "The principle that an event with very small probability will not happen is the only law of chance." (Borel, 1948)

While the Lewis-Loewer theory of objective chance is traditionally associated with the Principal Principle (Lewis even regarded PP as "non-negotiable"), it is very much – if not more – compatible with Cournot's principle: If the Humean probability of an event is very close to one, we can be almost certain that this event actually occurs. Why? Because this is what the best system is trying to tell us; because the way in which the Mentaculus summarizes relevant regularities in the mosaic is to assign to them a probability very close to one. Ironically, a version of what Lewis (1994) considered as the "big bad bug" of his theory of objective chance can serve to vindicate even the strongest form of CP in some cases. The Mentaculus will, for instance, assign a very small though positive probability to the universe evolving on an entropic-decreasing trajectory. However, if the universe actually did evolve on such an anti-entropic trajectory, this Mentaculus would not be the best systematization of our world (given that so many salient features depend on its entropic history). Hence, the fact that the Mentaculus assigns a near-zero probability to anti-entropic trajectories, together with the fact that the Mentaculus is the best systematozation of our world, implies that an entropy-decreasing evolution is, in fact, impossible. On the other hand, if we are talking about an event or regularity that the best system could conceivably fail to predict (probabilistically), it really doesn't matter if the Mentaculus assigns to it a chance of 10^{-50} or 10^{-100}. Our residual uncertainty about whetever that event obtains after all, does not come from anything the best system tells us about the world, but from the possibility that it just had to get this one wrong in the trade-off with simplicity.

In any case, so far, our main argument why even Humeans should favor CP over PP is that the concrete physical information that the Mentaculus provides is to be found, first and foremost, in statements of probability close to 1 and 0, while the rationality of aligning degrees of belief with any odd value of the Humean probability is spurious, at best. Note that also methodologically, the only way to *test* probabilistic laws is by some form of Cournot's principle, that is, by rejecting the law if we observe relevant phenomena to which it assigns a negligibly low probability (cf. Shafer and Vovk (2006)). I am not a verificationist and so do not claim that single-case probabilities are meaningless just because they cannot be empirically tested. I have, however, argued that the regularity account fails to give them meaning as deterministic chances – except to the extent that they can be reduced to typical frequencies. Humeans often claim that their probability measure is "empirical", yet provides information far beyond what is empirically testable. I don't think they can have it both ways.

4 Probability measures versus typicality measures

The next step from probability to typicality comes by emphasizing the following insight: If we agree that all we need are "probabilities" close to 1 and 0, then a whole lot of different measures could do the job. If we don't like the Lebesgue measure, how about putting a (truncated) Gaussian measure on M_{PH}? In fact, we can tweak the measure in almost any way we like. Almost any absolutely continuous measure will make a statement analogous to the FTSM true, and thus imply the same thermodynamic laws and statistical regularities. We cannot be too extreme, of course. A delta-measure concentrated on an anti-entropic microstate will, evidently, lead to very different predictions. However, as Maudlin (2007, p. 286) concludes: against this backdrop "our concerns about how to pick the 'right' probability measure to represent the possible initial states ... or even what the 'right' measure means, very nearly evaporate".

Something about these arguments may seem implausible. I have suggested that many measures agree with the normalized Liouville measure on probabilities close to 1 and 0, yet disagree significantly about the measure of other sets. Isn't the converse true, as well? Aren't there just as many measures that agree (more or less) on the probability of the United States electing a female president but disagree on the probability of entropy increase in the universe? Indeed, I have to be more precise about what I claim: Think of the (normalized) Liouville measure as a uniform probability density over the Past Hypothesis macroregion M_{PH}. If $\lambda(A) \approx 0$ but $\mu(A) \gg 0$, then μ must radically differ from λ on a small set A. (The same holds by contraposition for probabilities close to 1.) In contrast, for $\lambda(B) = 0.3$ while $\mu(B) = 0.4$ (let's say), μ needs to deviate only mildly from the uniform density over the relatively large set B. "Large" and "small" is here understood with respect to the Liouville measure, but this doesn't make the argument circular. The point is that radical deviations from the uniform Liouville measure would be necessary to come to different conclusions about typical/atypical events, while relatively small variations of the measure can lead to significantly different "probability" assignments for other events. More generally speaking, "probabilities" close to 1 and 0 are much more robust against variations of the measure than values in between.

For the sake of argument, Albert and Loewer are sometimes willing to concede that we could consider best system candidates that involve an entire set or equivalence class of probability measures, with the understanding that the theory endorses all and only those probability statements on which these measures (more or less) agree (cf. Albert (2015, footnote 2)). Of course, there is no good motivation to do so from a Humean perspective since a whole set or equivalence class of measures is neither simpler nor more informative than the Liouville measure (I have only argued that it is equally informative). The situation is rather the following: we should use the Liouville measure because it is simple and natural, though with the understanding that it is not a bona fide probability measure but a *typicality measure*. Its role and purpose is to designate certain events as "typical" or "atypical", while the numbers assigned to events that are neither (i.e., the values that are well between 0 or 1) come as mathematical surplus.

This move, to relate CP to a concept of typicality rather than probability, is a more recent development, though there is precedent for it in the physical literature (see e.g. Everett (1973) and the discussion in Barrett (2016), Bell (2004, Ch. 15), originally published in 1981, as well as Goldstein (2001) on Boltzmann). There are additional motivations for it:

1. As argued before, Cournot's principle suggests an understanding of probabilities as typical frequencies. As an account of (deterministic) probability, this would seem circular if "typical" were itself explicated in probabilistic terms. Moreover, probabilities can be understood as typical frequencies when applied to events *in* the universe, but we better not refer to frequencies (not even hypothetical ones) when we speak about the universe as a whole. For this reason alone, it makes sense to distinguish two different concepts.

2. In statistical mechanics, it is common and convenient to formalize typicality with the mathematical tools of measure theory – hence the deceptive kinship to probability. There are, however, other ways to define "typical", e.g., in terms of cardinalities or topological features, which can be relevant in other contexts and figure in the same way of reasoning. (For more subtle technical differences between typicality and probability, see Wilhelm (2019).)

3. Some authors have suggested that "typical" is a more intuitive and unambiguous notion than "probable" (see e.g. Dürr et al. (2017)). One way to spell this out is to say that the intuition associated with a typicality measure is one of "large" versus "small", rather than "probable" versus "improbable", sets (and that we have a better grasp of the former than the latter).

   Another way is to compare the following two formulations of Cournot's principle:

   i) Expect to find the regularities that obtain in nearly all possible worlds.

   ii) Expect to find the regularities that have a probability close to 1.

   I would argue that the first rationality principle has an immediate intuitive appeal. If there are facts about what obtains in nearly all nomologically possible worlds, it seems petty to ask why these facts are informative, or predictive, or explanatory. The second version seems more stipulative, or at least neutral with regard to the interpretative question, what fact the probability statement actually expresses and why it is supposed to guide our expectations. Moreover, ii) can be applied to any probability measure – and has meaningful content only in conjunction with a particular measure – while i) has intuitive appeal only to the extent that the measure on possible worlds captures the intuitive meaning of "nearly all".

5 The epistemic and metaphysical status of typicality

The last points are, admittedly, very controversial. Even some proponents of typicality are uncomfortable with the idea that there exists an a priori notion of "typical", and advocates of the Mentaculus emphatically deny that there are any typicality facts that come more or less for free once the rest of the theory is fixed. The deeper question here concerns the metaphysical status of the measure. According to the Mentaculus account, the probability measure has the same status as any other Humean law: it supervenes on the contingent regularities of the world as part of the best systematization. This option is, in principle, also available for the typicality measure; that is, one could consider a typicality- rather than a probability measure as part of the Humean best system (Callender, 2007). In my opinion, however, the Humean view – regardless of its flaws or merits as a metaphysics of laws, in general – fails to capture the more subtle aspects of typicality and its use in physics. These come across if one considers the typicality measure[3] not as another theoretical postulate on par with the dynamical laws but as a *way of reasoning about the dynamical laws*. In other words: being intimately tied to

[3] I am using the singular, though we should keep in mind that many different measures, qua mathematical objects, are equivalent as typicality measures.

Cournot's principle – which is normative rather than descriptive – the typicality measure falls itself, at least in part, into the normative domain.

Let me start to explore this by mapping out some key metaphysical differences between a typicality measure and a Humean probability measure:

1. According to the Mentaculus account, the Humean mosaic is the truthmaker of probability measure as part of the best system.

   According to the typicality account that I am defending, a choice of typicality measure can be *reasonable* or *justified*, but there are no concrete physical facts that make it, strictly speaking, true.[4]

2. According to the Mentaculus account, the probability measure together with the other best system laws cannot be entirely and radically wrong about the world, or else they would not form the best system.

   According to the typicality account, it is logically and metaphysically possible for a world to be – in any and all relevant regards – atypical with respect to the reference class of nomic possibilities.

3. Humean probabilities are supposed to summarize regularities in the actual world.

   Typicality statements summarize the modal structure of the laws. They do not refer directly to the actual world but to the fact that a certain feature or property is typical among all nomologically possible ones.[5]

All these points show that the typicality measure cannot be just another book-keeping device for the mosaic in the sense in which Humeans want to conceive of laws of nature in general. It can and should be tied to the dynamical laws (which, in turn, may be reducible or irreducible) but not in a way that depends substantially on contingent features of the world.

More precisely, I do not claim that the correct typicality measure follows analytically from the micro-dynamics, but I believe that there is no possible world in which the dynamical laws are the same as in ours while the notion of typicality is significantly different. Imagine, for instance, a world that is consistent with Newtonian dynamics but anti-entropic micro-conditions, a world, that is, in which apples spontaneously jump up in the air, in which gases tend to clump, and gravitating systems to be blown apart. I do not mean a world that is an exact time-reversal of ours (which would really be the

[4]Whether there are objective normative facts that make it true is beyond the scope of this paper.

[5]In other contexts, typicality may be understood with respect to a reference class that exists in the actual world; here, it has a decidedly modal character.

same world if there is no primitive direction of time) but one in which violations of our second law occur on a regular basis. It seems evident to me that the best system of this world – if one exists – would not involve Newtonian gravity with a strange typicality measure but very different microdynamics in the first place.

Or, if I can stop pretending to be Humean for a second: if the laws instantiated in that world were, in fact, the Newtonian ones, rational physicists would justifiable come to wrong conclusions about them. They would not discover the correct dynamical laws but have a radically different understanding of "typical" than we do. Typicality, at least when combined with an anti-Humean conception of laws, thus allows for the metaphysical possibility of "unlucky suckers" who find themselves in the absurd situation that rational inferences lead to radically wrong conclusions about the world they inhabit. Conversely, to the extent that the laws of nature are rationally accessible to us, our world must correspond – in the relevant respects – to a typical model of the true theory. This is nothing we could ever know with metaphysical certainty; we can only trust that "God is subtle but not malicious", as Einstein put it.

As so often, though, the epistemic situation for the Humean and anti-Humean is ultimately the same. While I think that it is always the theoretical system as a whole that is challenged by empirical evidence, I cannot conceive of a situation in which it seem rational to revise our notion of typicality instead of adjusting the dynamical laws (while the converse is common and unproblematic). For instance, there are almost certainly initial conditions for a Newtonian universe which are such that particles create an interference pattern whenever they are shot through a double-slit and recorded on a screen. This and other quantum phenomena are not made *impossible* by classical mechanics, they just come out as *atypical*. However, changing the typicality measure in order to give a Newtonian account of quantum phenomena is not a serious scientific option that anyone has ever, or should ever, entertain.

Using Quine's picture of a "web of belief" (Quine, 1951), I suggest that the dynamical laws are closer to the edges of the web than our notion of typicality. The typicality measure is somewhere in between the dynamical laws and the logical inference rules. This squares the circle between it being *necessary* but *not a priori*. While it is in principle possible to adjust it to new empirical evidence, the typicality measure is never the first knob to turn before making adjustments in other parts of the theoretical system.

One reason for this is that, because the notion of typicality is so robust against variations of the measure, any revision of it would be radical, i.e., must correspond to extreme changes in the measure on the space of possible micro-conditions. While it can be maintained that the change of the dynamics from Newtonian mechanics to quantum mechanics (or general relativity, etc.) was radical, as well, the new laws do at least recover the old ones in relevant limiting cases. It is hard to see how a similar continuity could hold between different standards of typicalilty. Instead, we would have to accept that we have been radically wrong about the meaning of "nearly all".

There is a more important reason why the typicality measure *should* be less empirical or epistemically more rigid than the dynamics. As mentioned in the introduction, due to the huge number of microscopic degrees of freedom, the dynamical laws put barely any constraints on what is physically possible on macroscopic scales. Given any macroscopic phenomenon, there are almost certainly microscopic initial conditions for which the laws of motion would account for it. By the same token, given almost any micro-dynamics and any phenomenon in the world, there will be some measure that makes the phenomenon "typical" or sufficiently likely, if you wish.

Therefore, treating the typicality (or probability) measure on the same footing as the dynamical laws would give us too many moving parts that can be adjusted to the data. For the Humean, this is bad because it increases the risk of a tie for the best system, at least in hypothetical situations when simplicity of the dynamical and probability postulates pull in opposite directions. For the anti-Humean, it is even worse, since the more freedom we have to adjust what counts as "typical" and "atypical", the less is empirical evidence able to inform us about what the true laws of our world actually are. Or, to put it differently: the more we regard the typicality measure as physically contingent, an independent empirical hypothesis, the less explanatory work is done by the dynamical laws and the class of possible worlds that they determine.

In fact, in actual scientific practice, typicality judgements *do* have a privileged status that the BSA fails to account for.[6] In particular, atypicality is precisely the standard by which dynamical theories are reasonably rejected as empirically inadequate (think again of the double-slit experiment as a falsification of classical mechanics or the 5σ-standard commonly used in particle physics). Interestingly, this applies in pretty much the same way to deterministic laws as to intrinsically stochastic ones. The difference is that in the latter case, falsifying a dynamical and a probabilistic hypothesis is one and the same, while in the deterministic case, it is primarily the dynamical laws that stand trial. And this is only possible because the typicality measure is epistemically more robust or, in some sense, entailed by the dynamics. In other words: in the scientific enterprise, some concept of "typicality" and "atypicality" is part of the backdrop against which law-hypotheses are evaluated rather than another law-hypothesis in its own right.

[6]Marc Lange (2009) makes a similar point when he argues for "degrees of necessity" in laws, though typicality is not quite a law and nomic necessity is not quite the right concept here.

6 Justification of the typicality measure

For these reasons, most advocates of typicality do not consider the typicality measure as an independent postulate of the physical theory, although it might be from a strictly logical point of view. But what then determines the right measure and accounts for its epistemic rigidity?

One answer we have already alluded to is that the role of the measure is not so much to *define* "typical" but to formalize an intuitive and largely pre-theoretic notion. In other words, there aren't competing versions of typicality corresponding to different choices of typicality measures, but one unified concept that a measure can either capture or fail to capture. If the set of possible worlds were finite, we would not feel the need for an additional postulate to express what we mean by "nearly all possible worlds". In the continuum case, there is more ambiguity about how to "count" but this is, arguably, a technical rather than a conceptual issue.

The concept of typicality has, in any case, a certain vagueness. Just as it is impossible, in general, to specify a precise threshold ratio above which we should say that a subset contains "nearly all" elements – i.e., a precise threshold measure close to 1 above which something counts as "typical" – it seems impossible to specify a fixed set of criteria that make a measure convincing as part of the mathematical formalization of typicality. In the context of classical Hamiltonian mechanics, the Liouville measure is clearly a reasonable choice while a delta-measure is clearly not, but a certain grey area in between seems unavoidable. Or consider of a family of Gaussian measures with standard deviation $\sigma \to 0$ so that the distributions become more and more peaked. Again, it would be misguided to ask for a sharp threshold value of σ below which the Gaussians cease to be suitable typicality measures. Still, this doesn't mean that the concept itself is ill-conceived or that the vagueness is problematic in practice. The bottom line is that a typicality statement has normative implications if and only if it is made with respect to a reasonable notion of "large" versus "small" sets. And while it is hard to state precisely what makes a measure reasonable or unreasonable, we can generally tell them apart when we see them.

Some authors put less emphasis on the intuitive content of typicality and more on the condition that the measure must be *stationary* under the dynamics. This is to say that the measure of a set of micro-states at one time must correspond to the measure of the time-evolved set at any other time. In this way, the dynamical laws would constrain the choice of typicality measure. I will provide a more precise definition of stationarity and further justification for this condition below.

Fortunately or unfortunately, as long as we are dealing with classical mechanics, both approaches are consistent and complementary, since the simplest and most intuitive measure – the uniform measure on phase space – is also stationary under the Hamiltonian dynamics (though not uniquely so). It could even be justified by a principle of indifference (Bricmont, 2001), though this is not a view that I will further entertain, mostly because I consider typicality facts to be independents of anyone's epistemic state.

There is, however, a notable example where stationarity and intuitiveness seem to go apart (and the principle of indifference to fail altogether). In Bohmian quantum mechanics, the natural typicality measure grounding Born's rule and thus quantum statistics for subsystems is given by the $|\Psi|^2$-density on configuration space, induced by the universal wave function Ψ (Dürr et al., 1992). This measure is stationary (more precisely: equivariant) under the Bohmian particle dynamics and even uniquely determined by this condition (Goldstein and Struyve, 2007). However, since we do not know what the universal wave function and hence the induced typicality measure looks like, it appears that its justification can hardly lie in pre-theoretic intuitions. If the $|\Psi|^2$-density turned out to be sharply peaked, it would look very different from a uniform measure and thus seem to constitute a radical departure from the notion of typicality employed in classical mechanics. In the next subsection, I will discuss how this conclusion can be avoided and the two approaches to the typicality measure reconciled.

6.1 Stationarity, uniformity, symmetry

The following discussion is quite technical, though the basic point is rather simple: So far, we have usually talked about measures on the initial macro-region of the universe. In fact, the relevant reference class of our typicality statements – what we actually want to "count" – are not initial conditions but (nomologically) possible worlds. Initial conditions are just a means to parametrize the respective solution trajectories. Stationarity, simply put, guarantees that a large set of solution trajectories is deemed large by the same standard (i.e. measure) if we look at the trajectories, i.e., the corresponding set of micro-states, at any other time.

Let $\mathcal{S}$ be the set of solution trajectories for the microscopic dynamics (consistent with the Past Hypothesis) in the state space $\Gamma \cong \mathbb{R}^n$. For any $t \in \mathbb{R}$, let $\varepsilon_t : \mathcal{S} \to \Gamma, X \mapsto X(t)$ be the map evaluating the trajectory X at time t. These maps can be read as charts, turning the solution set $\mathcal{S}$ into an n-dimensional differentiable manifold.[7] The transition maps between different charts are then $\varepsilon_t \circ \varepsilon_s^{-1} = \Phi_{t,s}$, where $\Phi_{t,s}$ is the flow arising as the general solution of the laws of motion (Fig. 2).

[7]Strictly speaking, some solutions may exist only on a finite time-interval so that the charts are only locally defined. Here, we assume global existence of solutions for simplicity.

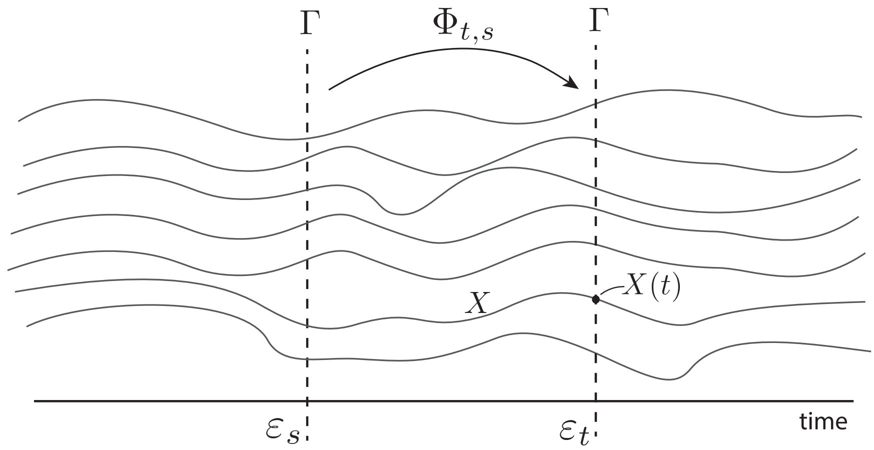

Figure 2: Sketch of the solution space and its parameterization by time slices.

Now, the easiest way to define a measure μ on S is in one of these charts, let's say ε_0. Indeed, a possible point of view is that there exists a distinguished "initial" time so that it makes sense to parametrize solutions by initial data at $t = 0$. The other point of view, emphasized by our geometric notation, is that the choice of the "time slice" is arbitrary, essentially amounting to a particular coordinatization of the solution space. Under a transition map (coordinate transformation), the measure transforms by a pullback, $\mu_t = \Phi_{t,0}\#\mu_0$, i.e. $\mu_t(A) = \mu_0(\Phi_{t,0}^{-1}A)$ for any measurable $A \subset \Gamma$, where μ_t is the measure represented in the chart ε_t.

A measure is *stationary* if and only if it has the same form in every time-chart, i.e.:

$$\mu_t = \Phi_{t,0}\#\mu_0 = \mu_0, \ \forall t \in \mathbb{R}, \tag{4}$$

Equivariance is the next best thing if the dynamics are themselves time-dependent. Concretely, in the case of Bohmian mechanics, the particle dynamics are determined by the universal wave function Ψ_t which may itself evolve in time according to a linear (Schrödinger) equation. Nonetheless, we have

$$|\Psi_t|^2 \mathrm{d}^n x = \Phi_{t,0}\# \left(|\Psi_0|^2 \mathrm{d}^n x\right), \ \forall t \in \mathbb{R} \tag{5}$$

so that the measure has the same functional form in terms of Ψ_t for any time t.

In conclusion, a stationary or equivariant measure on the state space Γ induces a canonical measure on the solution space $\mathcal{S}$; a measure that can be defined without distinguishing a set of coordinates, i.e., a particular moment in time.

Uniformity of a measure, on the other hand, is a metric notion. It requires that

$$\mu(B(x,r)) = \mu(B(y,r)),\ \forall x, y \in \Gamma, r > 0, \tag{6}$$

where $B(x,r)$ is the ball of radius r around x. However, even if the state space Γ comes equipped with a metric, it does not, in general, induce a canonical metric on $\mathcal{S}$. It is thus not clear what the metric on the solution space is supposed to be, or whether it makes sense to regard $\mathcal{S}$ as a metric space (Riemannian manifold) at all. But without a metric on the solution manifold, it is meaningless to ask whether a measure on it is uniform or not. From this point of view, it is indeed misleading to regard uniformity as a criterion for the typicality measure. Even the Liouville measure in classical mechanics is uniform on the "wrong space", namely on phase space rather than the space of possible worlds.

The uniformity of the Liouville measure on phase space does nonetheless capture a meaningful and important feature of the typicality measure, namely its *invariance under Galilean symmetries*. A symmetry is an isomorphism $T : \Gamma \to \Gamma$ that commutes with the flow, i.e. $\Phi_{t,s}(Tx) = T\Phi_{t,s}(x)$. This then induces a transformation $T^* : \mathcal{S} \to \mathcal{S}$ on the solution space by $T^* = \varepsilon_t^{-1} \circ T \circ \varepsilon_t$ which is independent of t.[8] The most important symmetries of classical mechanics are those of Galilean spacetime, namely:

$$\begin{aligned}
(q_1, ..., q_N; p_1, ..., p_N) \longrightarrow \quad & (q_1 + a, ..., q_N + a; p_1, ..., p_N) && \text{(Translation)}\\
& (Rq_1, ..., Rq_N; Rp_1, ..., Rp_N) && \text{(Rotation)}\\
& (q_1 + ut, ..., q_N + ut; p_1 + m_1 u, ..., p_N + m_N u) && \text{(Galilei boost)}
\end{aligned}$$

It is well-known that the Lebesgue or Liouville measure λ is invariant under these transformations (which are just Euclidean transformations on phase space). Consequently (as is easy to check), the induced measure on $\mathcal{S}$ is invariant under the corresponding symmetry transformations on the solution manifold.

In Bohmian mechanics, the issue is a bit more subtle since the wave function itself transforms non-trivially under Galilean symmetries, namely as (Dürr and Teufel, 2009):

$$\begin{aligned}
\Psi_t(q_1, ..., q_N) \longrightarrow \quad & \Psi_t(q_1 - a, ..., q_N - a) && \text{(Translation)}\\
& \Psi_t(R^{-1}q_1, ..., R^{-1}q_N) && \text{(Rotation)}\\
& \mathrm{e}^{\frac{i}{\hbar}\sum_{i=1}^N \left(uq_i - \frac{1}{2}m_i u^2\right)}\Psi_t(q_1 - ut, ..., q_N - ut) && \text{(Galilei boost)}
\end{aligned}$$

It is, however, evident that the $|\Psi|^2$-density is invariant under these transformations,

[8] Proof: $\varepsilon_t^{-1} T \varepsilon_t = \varepsilon_s^{-1} \Phi_{t,s} T \Phi_{s,t} \varepsilon_s = \varepsilon_s^{-1} \Phi_{t,s} \Phi_{s,t} T \varepsilon_s = \varepsilon_s^{-1} T \varepsilon_s$.

guaranteeing the invariance of the induced typicality measure under Galilean symmetries.

In the upshot, the typicality measures in both classical mechanics and Bohmian mechanics are justified and tied to the dynamics by precise mathematical features: stationarity/equivariance and Galilean invariance. However, at least in classical mechanics, these conditions are not sufficient to determine the measure uniquely or even rule out evidently inadequate choices (such as a delta-measure concentrated on a stationary microstate). At the end of the day, part of what makes a measure compelling and allows it to play its normative role as it figures in typicality reasoning is that the choice doesn't seem biased or ad hoc or overly contrived. In other words, I do not believe that "soft" criteria can or should be completely avoided, and they aren't, in fact, in scientific practice. Attempts to axiomatize typicality measures (Werndl, 2013) seem misguided, not just because the particular proposals are uncompelling but because it is hard to see why any set of formal axioms should be more compelling than the very measures we use in physics.

7 The lottery paradox and rational belief

We have said that (Humean) probabilities are usually tied to the Principal Principle while typicality is tied to Cournot's principle. The Principal Principle is formulated in doxastic terms, that is, in terms of credences or *degrees of belief*. If we wanted to formulate Cournot's principle correspondingly – which some advocates of typicality would reject – it would be tied to *belief simpliciter*. One problem that then arises, is that rational belief is usually assumed to be closed under conjunction ($Bel(A) \wedge Bel(B) \implies Bel(A \wedge B)$) while typicality statements are not[9]. Clearly, the measure of $A_1 \cap A_2 \cap \ldots \cap A_n$ (reading A_i is the set of possible worlds in which A_i is true) could get arbitrarily small with increasing n, even if the measure of each A_i is very large.

The account thus runs into the infamous *lottery paradox* (Kyburg, 1961): It is typical for any lottery ticket to be a loser. Hence, I should believe that ticket 1 will lose, and that ticket 2 will lose, ..., and that ticket N will lose. But then I must also believe "ticket 1 will lose and ticket 2 will lose ... and ticket N will lose", which is certainly false if one of the tickets will be drawn for sure.

A possible reaction is to bite the bullet and admit that belief (or at least the relevant notion of "expectation" associated with typicality) is not closed under conjunction. After all, for anyone participating in the lottery, it is wise to act as if her ticket is going to lose (rather than buying a boat right away), while the lottery company better act as

[9]Unless one considers a very strong notion of typicality defined in terms of sets of measure 1.

if they will have to pay out the jackpot eventually. Moreover, the following variant of the lottery paradox seems plausibly true: it is typical that *some* atypical events actually occur.

A more sophisticated response was developed by Hannes Leitgeb (2014, 2017) who tackled the problem of reconciling normal doxastic logic and degrees of belief (satisfying the axioms of probability) with the "Lockean thesis": there exists a threshold $\frac{1}{2} < r < 1$ such that any proposition A is believed if and only if the degree of belief in A is at least r. Formally: $Bel(A) \iff C(A) \geq r$. In a nutshell, Leitgeb proves that this can be done for the price of admitting that belief is *context sensitive*, i.e., that the threshold value r, the credence function C, and the relevant set of propositions Π can be codependent. In his *stability theory of belief*, there then exists, in any given context, a unique proposition B_W of stably high subjective probability ($C(B_W \mid A) > \frac{1}{2}$ for all compatible propositions A) such that $Bel(B) \iff B_W \subseteq B$.

The parallels to typicality and the Cournot principle are evident, though it will require more work to lay them out in detail. In any case, while I have argued that typicality is a uniform concept, what we consider to be "typical" or "atypical" is indeed context-dependent, and the threshold seems to be different, in an interesting hierarchical way, for regularities falling under the purview of fundamental physics or the various special sciences. Adapting Leitgeb's stability theory: In our discussion of macroscopic "laws", the fundamental theorem of statistical mechanics may be the basic typicality fact that then grounds or implies the typicality of thermodynamic and statistical regularities; but me winning the lottery is not the kind of event that we judge in the same context. Which is also why its occurrence would be extremely surprising to me and my peers but not to the extent that it would put our fundamental theories of physics into question.

As a side note: Emile Borel (1939), discussing Cournot's principle (though not by this name), actually made a proposal for the orders of magnitude characterizing "negligible probabilities": $< 10^{-6}$ on the individual human scale, $< 10^{-15}$ on the terrestrial scale, $< 10^{-50}$ on the cosmological scale. He argued: "In the ordinary conduct of his life, every man usually neglects probabilities whose order of magnitude is less than 10^{-6}, that is, one millionth, and we will even find that a man who would constantly take such unlikely possibilities into account would quickly become a maniac or even a madman." The spirit of Borel's reflections is more interesting here than the precise numbers.

8 Conclusion

Our discussion shows that the debate between the Mentaculus and the typicality account involves at least three separate issues:

1. Is the measure grounding the predictions of statistical mechanics a bona fide probability measure or a typicality measure?

2. What is the epistemic and metaphysical status of the measure? Is it a theoretical postulate (a Humean law) on par with the dynamics, or does it formalize a way of reasoning about the laws?

3. What expresses its epistemic or behavior-guiding function, the Principal Principle or Cournot's principle?

While I have defended a view that comes down on the opposite side of the Mentaculus in every relevant respect, it is important to note that the answers to these three questions are logically independent – with the exception that a typicality measure essentially collapses the Principal Principle into Cournot's principle. This leaves much room for compromise and moderation, but also for misunderstandings, since not everyone defending "typicality" or "Cournot's principle" may have the same package deal in mind.

It is possible to maintain, for instance, that the (Humean) laws involve a typicality measure rather than a probability measure, that a probability measure expresses a way of reasoning (e.g., a principle of indifference) rather than an empirical postulate, or that the laws of nature include a bona fide probability measure in addition to deterministic laws, although only probabilities near 1 and 0 are physicall relevant.

Of course, while there is always value in compromise and moderation, the extremal positions are often the most interesting ones. Here, one of them, the Mentaculus theory, is attractive because of its stringency, cohesiveness and philosophical potential, though I have argued that it is both too ambitious and simplistic in its treatment of single-case probabilities and fails to capture the different (epistemic) status that dynamical postulates and probabilistic judgements have in physics. The opposite view – which I have defended as the typicality account – seems more compelling as an account of good scientific reasoning and of how macroscopic regularities are grounded in microscopic laws, though much of its potential remains still untapped.

References

Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.

Albert, D. Z. (2015). *After Physics*. Cambridge, Massachusetts: Harvard University Press.

Barrett, J. A. (2016). Typicality in Pure Wave Mechanics. *Fluctuation and Noise Letters*, 15(03):1640009.

Bell, J. S. (2004). *Speakable and Unspeakable in Quantum Mechanics*. Cambridge: Cambridge University Press, second edition.

Borel, E. (1939). *Valeur pratique et philosophie des probabilités*. Gauthier-Villars.

Borel, É. (1948). *Le hasard*. Presses universitaires de France.

Bricmont, J. (1995). Science of Chaos or Chaos in Science? *Annals of the New York Academy of Sciences*, 775(1):131–175.

Bricmont, J. (2001). Bayes, Boltzmann and Bohm: Probabilities in Physics. In Bricmont, J., Ghirardi, G., Dürr, D., Petruccione, F., Galavotti, M. C., and Zanghi, N., editors, *Chance in Physics: Foundations and Perspectives*, pages 3–21. Berlin: Springer.

Callender, C. (2007). The emergence and interpretation of probability in Bohmian mechanics. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(2):351–370.

Carroll, S. (2010). *From Eternity to Here*. Dutton, New York.

Cournot, A. A. (1843). *Exposition de la théorie des chances et des probabilités*. L. Hachette.

Dürr, D., Froemel, A., and Kolb, M. (2017). *Einführung in Die Wahrscheinlichkeitstheorie Als Theorie Der Typizität*. Springer Berlin Heidelberg, Berlin, Heidelberg.

Dürr, D., Goldstein, S., and Zanghí, N. (1992). Quantum equilibrium and the origin of absolute uncertainty. *Journal of Statistical Physics*, 67(5-6):843–907.

Dürr, D. and Teufel, S. (2009). *Bohmian Mechanics: The Physics and Mathematics of Quantum Theory*. Berlin: Springer.

Everett, H. (1973). The Theory of the Universal Wave Function. *The Many-Worlds Interpretation of Quantum Mechanics*, pages 3–140.

Goldstein, S. (2001). Boltzmann's Approach to Statistical Mechanics. In Bricmont, J., Dürr, D., Galavotti, M. C., Ghirardi, G., Petruccione, F., and Zanghì, N., editors, *Chance in Physics: Foundations and Perspectives*, pages 39–54. Springer, Berlin.

Goldstein, S. (2012). Typicality and Notions of Probability in Physics. In *Probability in Physics*, The Frontiers Collection, pages 59–71. Springer, Berlin, Heidelberg.

Goldstein, S. and Struyve, W. (2007). On the uniqueness of quantum equilibrium in Bohmian mechanics. *Journal of Statistical Physics*, 128(5):1197–1209.

Hájek, A. (2006). The reference class problem is your problem too. *Synthese*, 156:563–585.

Hall, N. (1994). Correcting The Guide to Objective Chance. *Mind*, 103(412):505–518.

Hall, N. (2004). Two Mistakes About Credence and Chance. *Australasian Journal of Philosophy*, 82(1):93–111.

Kolomogoroff, A. (1933). *Grundbegriffe Der Wahrscheinlichkeitsrechnung.* Ergebnisse Der Mathematik Und Ihrer Grenzgebiete. 1. Folge. Springer-Verlag, Berlin Heidelberg.

Kyburg, H. E. (1961). *Probability and the Logic of Rational Belief.* Wesleyan University Press.

Lange, M. (2009). *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature.* Oxford University Press, Oxford, New York.

Lazarovici, D. and Reichert, P. (2015). Typicality, Irreversibility and the Status of Macroscopic Laws. *Erkenntnis*, 80(4):689–716.

Leitgeb, H. (2014). The Stability Theory of Belief. *Philosophical Review*, 123(2):131–171.

Leitgeb, H. (2017). *The Stability of Belief: How Rational Belief Coheres with Probability.* Oxford University Press.

Lewis, D. (1980). A subjectivist's guide to objective chance. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability. Volume II*, pages Reprinted in David Lewis (1986): Philosophical papers. Volume 2. Oxford: Oxford University Press. Pp. 83–132. Berkeley: University of California Press.

Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103(412):473–490. Reprinted in Lewis, D. (1999). \emphPapers in Metaphysics and Epistemology, pages 224-247. Cambridge: Cambridge University Press.

Loewer, B. (2001). Determinism and Chance. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 32(4):609–620.

Loewer, B. (2004). David Lewis's Humean Theory of Objective Chance. *Philosophy of Science*, 71(5):1115–1125.

Loewer, B. (2007). Counterfactuals and the Second Law. In Price, H. and Corry, R., editors, *Causation, Physics, and the Constitution of Reality. Russell's Republic Revisited*, pages 293–326. Oxford: Oxford University Press.

Loewer, B. (2012a). The emergence of time's arrows and special science laws from physics. *Interface Focus*, 2(1):13–19.

Loewer, B. (2012b). Two accounts of laws and time. *Philosophical Studies*, 160(1):115–137.

Martin, T. (1996). *Probabilités et critique philosophique selon Cournot*. Vrin.

Maudlin, T. (2007). What could be objective about probabilities? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(2):275–291.

Oldofredi, A., Lazarovici, D., Deckert, D.-A., and Esfeld, M. (2016). From the Universe to Subsystems: Why Quantum Mechanics Appears More Stochastic than Classical Mechanics. *Fluctuation and Noise Letters*, 15(03):1640002.

Penrose, R. (1999). *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*. OUP Oxford, new edition edition.

Quine, W. V. (1951). Main Trends in Recent Philosophy: Two Dogmas of Empiricism. *The Philosophical Review*, 60(1):20–43.

Schwarz, W. (2014). Proving the Principal Principle. In Wilson, A., editor, *Chance and Temporal Asymmetry*, pages 81–99. Oxford University Press.

Shafer, G. and Vovk, V. (2006). The Sources of Kolmogorov's Grundbegriffe. *Statistical Science*, 21(1):70–98.

Volchan, S. B. (2007). Probability as typicality. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(4):801–814.

Werndl, C. (2013). Justifying typicality measures of Boltzmannian statistical mechanics and dynamical systems. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(4):470–479.

Wilhelm, I. (2019). Typical: A Theory of Typicality and Typicality Explanation. *The British Journal for the Philosophy of Science*.