

Lingnan University

Digital Commons @ Lingnan University

Lingnan Theses and Dissertations

Theses and Dissertations

10-12-2020

Worlds in a stochastic universe : on the emergence of world histories in minimal Bohmian Mechanics

Alexander EHMANN

Follow this and additional works at: <https://commons.ln.edu.hk/otd>



Part of the [Metaphysics Commons](#)

Recommended Citation

Ehmann, A. (2020). Worlds in a stochastic universe: On the emergence of world histories in minimal Bohmian Mechanics (Doctor's thesis, Lingnan University, Hong Kong). Retrieved from <https://commons.ln.edu.hk/otd/81/>

This Thesis is brought to you for free and open access by the Theses and Dissertations at Digital Commons @ Lingnan University. It has been accepted for inclusion in Lingnan Theses and Dissertations by an authorized administrator of Digital Commons @ Lingnan University.

Terms of Use

The copyright of this thesis is owned by its author. Any reproduction, adaptation, distribution or dissemination of this thesis without express authorization is strictly prohibited.

All rights reserved.

WORLDS IN A STOCHASTIC UNIVERSE:
ON THE EMERGENCE OF WORLD HISTORIES
IN MINIMAL BOHMIAN MECHANICS

EHMANN ALEXANDER

PHD

LINGNAN UNIVERSITY

2020

WORLDS IN A STOCHASTIC UNIVERSE:
ON THE EMERGENCE OF WORLD HISTORIES
IN MINIMAL BOHMIAN MECHANICS

by
EHMANN Alexander

A thesis
submitted in partial fulfillment
of the requirements for the Degree of
Doctor of Philosophy in Philosophy

Lingnan University

2020

ABSTRACT

Worlds in a Stochastic Universe: On the Emergence of World Histories in Minimal Bohmian Mechanics

by

EHMANN Alexander

Doctor of Philosophy

This thesis develops a detailed account of the emergence of for all practical purposes continuous, quasi-classical world histories from the discontinuous, stochastic micro dynamics of Minimal Bohmian Mechanics (MBM). MBM is a non-relativistic quantum theory. It results from excising the guiding equation from standard Bohmian Mechanics (BM) and reinterpreting the quantum equilibrium hypothesis as a stochastic guidance law for the random actualization of configurations of Bohmian particles. On MBM, there are no continuous trajectories linking up individual configurations. Instead, individual configurations are actualized independently of each other, carving out the decoherence-induced branching structure of the universal wave function. Yet, by contrast to the Everett interpretation, branches of the universal wave function are not actualized in parallel, i.e. all at the same time. Rather, world branches, on MBM, are actualized sequentially.

For an introduction to MBM, the transition from BM to MBM is described, and their empirical equivalence is established. I present the conditions under which MBM can be classified as a primitive ontology (PO) theory, regarded as crucial for the acceptance of a theory as Bohmian by many proponents of BM. While rendering MBM compatible with the PO approach, it must not fall prey to the problem of communication, arising from the two-space reading of BM.

The issue of temporal solipsism, identified by Bell as a serious problem for his “Everett (?) theory” – the historic predecessor of MBM – is discussed. I argue that Barbour’s time capsule approach does not provide a satisfactory solution. In particular, his adopting a more than minimal psychophysical parallelism between brain processes and experience of macroscopic change is argued to be reasonable in light of our current best neuroscience, yet problematic for theories relying solely on time capsules, understood as highly structured, individual, internally static configurations. As a solution, I introduce worlds, and world histories, as key concepts in providing a link between MBM’s micro dynamics and macroscopic phenomena. Worlds, other than time capsules, are coarse-grained regions in phase- or configuration space, defined as sets of possible micro states satisfying a relation of sufficient similarity from a macroscopic perspective, with respect to a given micro state. Hence, worlds may overlap.

I argue that worlds thus construed are a reasonable option for replacing the disjoint macro regions of phase space, resulting from the usual way of partitioning phase space in the standard framework of Boltzmannian statistical mechanics. This move solves the issue of discontinuous change of macro variables upon the micro state crossing the boundary between different macro regions.

I adapt this move for MBM's discontinuous micro dynamics in configuration space. Issues revolving around the microscopic-macroscopic distinction, particle identity, impenetrability and haecceity in light of the desideratum of particle number conservation, etc., are discussed. I provide a detailed explanation of how overlapping worlds in MBM form world histories, thereby linking up macroscopically distinct worlds. Thus, the problem of temporal solipsism is resolved in a way that is compatible with a more than minimal psychophysical parallelism.

DECLARATION

I declare that this is an original work based primarily on my own research, and I warrant that all citations of previous research, published or unpublished, have been duly acknowledged.

A rectangular box with diagonal hatching containing a handwritten signature in black ink. The signature is stylized and appears to be 'EHMANN Alexander'.

SIGNED

(EHMANN Alexander)

CERTIFICATE OF APPROVAL OF THESIS

WORLDS IN A STOCHASTIC UNIVERSE:
ON THE EMERGENCE OF WORLD HISTORIES
IN MINIMAL BOHMIAN MECHANICS

by

EHMANN Alexander

Doctor of Philosophy

Panel of Examiners :

SIGNED

(Chairman)

(Prof. HAMPTON Mark Andrew)

SIGNED

(External Member)

(Prof. CALLENDER Craig A.)

SIGNED

(Internal Member)

(Prof. ROWBOTTOM Darrell Patrick)

SIGNED

(Internal Member)

(Prof. MARSHALL Daniel Graham)

Chief Supervisor :

Prof. ROWBOTTOM Darrell Patrick

Co-supervisor :

Prof. ZHANG Jiji

Approved for the Senate :

SIGNED

(Prof. MOK Ka Ho Joshua)

Chairman, Postgraduate Studies Committee

12 OCT 2020

Date

CONTENTS

ACKNOWLEDGEMENTS.....	iii
PRELUDE: ZENO'S ARROW	1
I. INTRODUCTION AND OVERVIEW	3
II. A SHORT INTRODUCTION TO MINIMAL BOHMIAN MECHANICS	7
II.1 BOHMIAN MECHANICS AS A SOLUTION TO THE MEASUREMENT PROBLEM	10
II.2 A GENERAL DESCRIPTION OF BM AND ITS EMPIRICAL EQUIVALENCE WITH QM	13
II.3 FROM BOHMIAN MECHANICS TO MINIMAL BOHMIAN MECHANICS.....	15
II.4 MBM AS A PO THEORY AND THE PROBLEM OF COMMUNICATION	20
<i>II.4.1 What makes a PO theory?</i>	<i>20</i>
<i>II.4.2 Is MBM compatible with the PO approach?</i>	<i>24</i>
<i>II.4.3 The problem of communication</i>	<i>25</i>
<i>II.4.4 Disposition or law?</i>	<i>27</i>
II.5 A BIT OF HISTORY: BELL'S EVERETT (?) THEORY AND TEMPORAL SOLIPSISM.....	30
II.6 BARBOUR: TEMPORAL SOLIPSISM? TIME CAPSULES!	33
<i>II.6.1 Explaining experiences.....</i>	<i>35</i>
<i>II.6.2 Fine-tuning for life and rare time capsules: an analogy</i>	<i>39</i>
<i>II.6.3 Explaining the need for explaining experiences: a conjecture</i>	<i>41</i>
<i>II.6.4 Is MBM empirically incoherent?</i>	<i>48</i>
II.7 TEMPORAL SOLIPSISM SOLVED?	49
III. SMOOTHING OUT BOLTZMANNIAN STATISTICAL MECHANICS: ON THE VIABILITY OF AN ALTERNATIVE PARTITIONING OF PHASE SPACE INTO OVERLAPPING WORLDS	57
III.1 THE USUAL PRESENTATION OF BSM	60
III.2 REMARKS ON MACRO REGIONS AND PARTITIONING PHASE SPACE.....	61
III.3 THE DISCONTINUITY PROBLEM AS IT ARISES FROM THE FRAMEWORK OF BSM	64
III.4 BSM AS A FUNDAMENTAL THEORY IN LIGHT OF THE DISCONTINUITY PROBLEM	71
III.5 HINTS OF DIGRESSION FROM THE USUAL.....	73
III.6 AGAINST SHARP BOUNDARIES OF MACRO REGIONS IN PHASE SPACE (I): CROSSING BOUNDARIES	75
III.7 AGAINST SHARP BOUNDARIES OF MACRO REGIONS IN PHASE SPACE (II): SHIFTING BOUNDARIES	77
III.8 AGAINST SHARP BOUNDARIES OF MACRO REGIONS IN PHASE SPACE (III): ENTROPY	81
III.9 TOWARDS AN ALTERNATIVE	82
<i>III.9.1 Removing exhaustive-cover.....</i>	<i>83</i>
<i>III.9.2 Removing no-overlap.....</i>	<i>85</i>

III.9.3 <i>A different solution?</i>	87
III.10 FROM MACRO REGIONS TO WORLDS	87
III.10.1 <i>Introducing a new framework for BSM: Some conditions</i>	89
III.10.2 <i>Back to the example</i>	100
III.10.3 <i>World size, the equilibrium world, and entropy increase</i>	101
III.11 CONCLUSION	104
IV. WORLDS IN A STOCHASTIC UNIVERSE.....	106
IV.1 WORLDS FROM THE WAVE FUNCTION VS. WORLDS FROM THE PARTICLE CONFIGURATION.....	107
IV.1.1 <i>Many worlds in denial?</i>	108
IV.1.2 <i>Psychophysical Parallelism</i>	112
IV.2 WHAT IS A WORLD?.....	114
IV.3 SUPERPOSITIONS IN MBM.....	120
IV.4 THE MICROSCOPIC-MACROSCOPIC-DISTINCTION AND ITS IMPORT ON THE CONSTRUAL OF WORLDS.....	125
IV.4.1 <i>The Microscopic and the Macroscopic Domain vs. the Quantum and the Classical Domain</i>	129
IV.5 A MINOR ISSUE WITH BARBOUR'S ACCOUNT	131
IV.6 THE MACROSCOPIC OBJECT.....	138
IV.7 CONFIGURATIONS AND CONFIGURATION SPACES.....	143
IV.7.1 <i>Particle identity, impenetrability, and the desideratum of particle number conservation in MBM</i>	147
IV.7.2 <i>Particle haecceity vs. haecceitism</i>	150
IV.7.3 <i>Individuation of particles qua their space-time trajectories</i>	152
IV.8 WHAT IS A WORLD? REDUX	155
IV.9 FROM DISCRETE CONFIGURATIONS TO CONTINUOUS WORLD HISTORIES VIA OVERLAPPING WORLDS.....	157
IV.9.1 <i>Relative frequency of world histories?</i>	163
IV.9.2 <i>Quasi-parallel world histories</i>	165
IV.9.3 <i>World histories: Splitting and merging</i>	168
IV.9.4 <i>An application of splitting world histories: The double-slit experiment</i>	172
IV.9.5 <i>Counting world histories</i>	177
IV.10 MBM IS COMPATIBLE WITH DOUBLE ASPECT PPP.....	178
V. CONCLUSION AND OUTLOOK	179
REFERENCES	183

ACKNOWLEDGEMENTS

I gratefully acknowledge generous funding of my research at Lingnan University by the Research Grants Committee of the University Grants Council of Hong Kong through the Hong Hong PhD Fellowship Scheme.

I gratefully acknowledge generous funding of my stay as a visiting PhD student at the University of Salzburg by the Austrian Federal Ministry of Education, Science and Research through the Ernst Mach Grant (worldwide).

I am indebted to my supervisors Darrell Rowbottom and Jiji Zhang (both Lingnan) for their invaluable feedback and their continuous support.

I am grateful to Charlotte Werndl (Salzburg, LSE) and Patricia Palacios (Salzburg, MCMP) for their supervision of my research at the University of Salzburg and their support of my application to the Ernst Mach Grant.

I thank Craig Callender (UCSD) for being the external examiner of my thesis, Darrell Rowbottom (Lingnan) and Dan Marshall (Lingnan) for being the internal examiners of my thesis, and Mark Hampton (Lingnan) for being the chairman of my oral examination.

Vinnie Cheung and Catherine Cheung deserve special mention. They are the heart and soul of Lingnan's Department of Philosophy.

Last but not least, I thank my dear friend and collaborator Patrick Dürr (Oxford) for countless hours of discussion.

I dearly hope that Hong Kong and its great universities will remain places where unconventional, even controversial thought is not only developed, but expressed and discussed freely and peacefully, as is essential for the human endeavour that is the progress of science, and for the endeavour that is the progress of humanity.

Prelude: Zeno's arrow

*For a moment, nothing happened.
Then, after a second or so, nothing continued to happen.*
— Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

In a moment, nothing happens. In particular, nothing moves. Zeno's arrow is testament to that: in any given moment, it doesn't move. So it moves never. It doesn't move at all.

And yet, arrows move. At least so it seems.

In any given point, a circle has a unique tangent. A point by itself has no tangent. What's the difference between a point, and a point on a circle? A point is a point, one should think, and rightly so. Why does a point on a circle have a tangent, but not a point on its own? That's to mistake one thing for another. It's not the point on the circle that has a unique tangent. It's the circle that has a unique tangent in that point. Hold that point fixed and take away the circle. The tangent disappears.

The velocity of the arrow in any given moment is like the tangent of the circle in any given point. Hold that moment fixed and take away the other moments. The velocity disappears.

The circle consists of a continuum of points, including the tangential point. Take one of them, left or right of the tangential point, and draw a straight line through them. Let the new point approach the tangential point. In the limit, the line is the tangent of the circle in that point. You can't construct a tangent like this from a single point,

because a single point isn't part of a continuum of points on which you could approach that single point.

Take two moments of the arrow's flight. They are separated by a finite time interval. Look at the positions of the arrow at these moments. Between the two, the arrow was displaced, by a certain amount, in a certain direction. Divide this displacement by the length of the time interval, and you get the average velocity over that time interval. Let the time interval approach zero. The limit is the instantaneous velocity of the arrow at that moment in time. You can't have an instantaneous velocity like this from a single moment, because a single moment isn't part of a continuum of moments on which you could approach that single moment.

Crudely put, one of the basic posits of Bohmian Mechanics says that particles, despite behaving non-classically, at any moment have an instantaneous velocity. They move on continuous, deterministic trajectories, like an arrow. One of the basic posits of Minimal Bohmian Mechanics contests this. It says that particles don't have an instantaneous velocity, ever. There are no continuous, deterministic trajectories. This is the core difference between Bohmian Mechanics and Minimal Bohmian Mechanics.

This thesis is about Minimal Bohmian Mechanics and the claim that arrows can move despite particles having no instantaneous velocity. The matter is somewhat intricate, as you will see, but the overall insight is this: although, according to Minimal Bohmian Mechanics, things don't move on continuous trajectories on the fundamental level of description, on a higher, emergent level of description, they exhibit continuous motion. This is made possible by what I call the emergence of for all practical purposes continuous, quasi-classical world histories from the discontinuous, stochastic micro dynamics of particles.

Similar accounts have been devised before. Most famously, John Bell discussed a predecessor of Minimal Bohmian Mechanics, but immediately rejected it, for reasons to be made explicit below. I hold that his dismissal was premature. Almost as famously, Julian Barbour has developed a fundamentally timeless cosmology, quite similar in spirit to Bell's theory. Although Barbour doesn't immediately reject his own theory, I hold that both, Bell's and Barbour's account, are deficient in some respects. With my presentation of emergent world histories in Minimal Bohmian Mechanics, I propose a way to overcome these deficiencies.

Let's make the arrow fly.

I. Introduction and overview

This thesis develops a detailed account of the emergence of for all practical purposes (FAPP) continuous, quasi-classical world histories, containing FAPP-continuous macroscopic phenomena, from the discontinuous, stochastic micro dynamics of Minimal Bohmian Mechanics (MBM). MBM is a non-relativistic quantum theory, recently put forward by Patrick Dürr and me. It results from excising the guiding equation from standard Bohmian Mechanics (BM) and reinterpreting the quantum equilibrium hypothesis as a stochastic guidance law for the random actualization of configurations of Bohmian particles. On MBM, there are no continuous trajectories linking up individual configurations. Instead, individual configurations are actualized independently of each other, carving out the decoherence-induced branching structure of the universal wave function. Yet, by contrast to the Everett interpretation, branches of the universal wave function are not actualized in parallel, i.e. all at the same time. Rather, world branches, on MBM, are actualized sequentially.

Along the way of developing my approach of emergent world histories, I will be discussing several, interrelated topics. First, I will give a short introduction to MBM

(II), using standard BM as the background against which it stands out. I will discuss BM's solution of the eminent measurement problem (II.1) and its empirical equivalence with orthodox QM (II.2). The transition from standard BM to MBM will be made by excising the guiding equation from the former and reinterpreting the quantum equilibrium hypothesis as a stochastic law, governing MBM's fundamentally random, discontinuous particle dynamics (II.3). I will argue that MBM is empirically equivalent with standard BM, and thus with QM, despite its lack of continuous, deterministic particle trajectories. Since many proponents of BM take the latter to be a primitive ontology (PO) theory, and regard this as somewhat crucial, I will argue that MBM can be made compatible with the PO approach, given certain restrictions on the interpretation of its WF (II.4). MBM to some extent resembles Bell's "Everett (?) theory". The latter was immediately rejected by its creator on the charge of leading to the problem of temporal solipsism (II.5). I will locate the root of this problem in Bell's identification of instantaneous configurations with worlds: on his construal, each world contains one, and only one, configuration. Accordingly, each world is likely to be instantiated only once. Barbour's time-capsule approach, at first sight, suggests itself as a possible solution to the problem of temporal solipsism. Time capsules are a special type of configurations: they contain consistent records of past events that, according to Barbour, suffice for the reconstruction of consistent histories from within individual, static configurations. However, Barbour's time-capsule approach provides no account of how temporally remote configurations are linked. Like Bell, Barbour identifies individual configurations – albeit special ones – with worlds. Thus, his approach isn't suitable to solve the problem of temporal solipsism within the MBM framework (II.6).

In order to tackle this issue, I will propose to adopt a notion of worlds different from the one put forward by Bell and Barbour. To introduce this notion, I am going to discuss the viability of my approach for solving what I call the discontinuity problem of Boltzmannian statistical mechanics (BSM) in (III). Following the usual presentation of the framework (III.1), the available phase space of a system is

partitioned into a finite set of disjoint macro regions, corresponding to different macro states (III.2). This construal of macro regions leads to the discontinuity problem of BSM: upon the micro state evolving from one macro region to another, it crosses boundaries between them, thereby instantiating different macro values associated with these macro regions. As a result, the usual framework of BSM invariably models change of macro values, whenever they change at all, as a jump discontinuity (III.3). I argue that this is especially problematic for those who take BSM to be a fundamental theory (III.4). The discontinuity problem has not been made explicit in the literature so far, although there are some hints to be found, to the effect that the usual way of partitioning phase space into disjoint sets with sharp boundaries is not considered entirely unproblematic by the community (III.5). Hence, I will present its consequences in a series of examples (III.6, III.7, III.8). I will consider two seemingly obvious approaches to tackle the problem, removing the underlying assumptions of the standard framework. As we will see, none leads to acceptable results (III.9). Thus, I will introduce a new framework for BSM (III.10) that allows carving up phase space, not into disjoint macro regions, but potentially overlapping worlds, defined as sets of phase points sufficiently similar to a given micro state, i.e. instantiating macro states indistinguishable from a given macro state. Worlds thus construed, other than macro regions, can overlap consistently. As a result, the need to model the change of macro values as discontinuous jumps can be avoided. I argue that the spirit and basic features of BSM remain intact despite replacing macro regions with worlds. In particular, worlds, like macro regions, can be compared with other worlds (and the entirety of available phase space) in terms of their volumes, thus allowing for similar accounts of equilibrium states and approach to equilibrium as usual.

In the next chapter (IV), I will make use of this notion of worlds to solve a similar discontinuity problem as it arises in MBM. There, it is put into sharp relief: not only do remote configurations instantiate different macro states, but individual configurations aren't linked up by continuous trajectories. That is, not even on the level of the microscopic description must a continuous evolution be assumed. This

lack of continuity in the micro evolution of the universe, following Bell, leads to the problem of temporal solipsism. To circumvent it, one could simply adopt the Everettian notion of worlds, understood as entire, decoherence-induced world branches. In particular, if MBM should turn out to be many worlds in denial, such that the particle configuration turns out to be superfluous theoretical structure – as, according to Brown, is the case in BM – this move would be appropriate. However, MBM features fundamental particles. At least in this respect, it is a bona fide Bohmian theory. In Bohmian theories, the natural basis for construing worlds are particle configurations, not the WF (IV.1). Also, Brown’s criticism of standard BM doesn’t carry over to MBM. The latter is not many worlds in denial; its particles are not superfluous theoretical structure. Rather, they are the substance material stuff is made of (IV.1.1). With particles playing the role of a substance in MBM’s ontology, an account of psychophysical parallelism should take them to be the physical aspect of the psychophysical. Thus, worlds should be construed from particle configurations (IV.1.2). Accordingly, I will introduce the notion of worlds as sets of particle configurations into the framework of MBM (IV.2). For this, I make use of a nontransitive relation – sufficient similarity – already familiar from the previous chapter on BSM. In order to make my approach work, I discuss the issue of superpositions in MBM (IV.3), and the import of the microscopic-macroscopic distinction on the construal of worlds (IV.4). I will argue that Barbour’s time-capsule approach is not suitable for the purpose of construing worlds, and ultimately, world histories (IV.5). It is deficient, in particular, when taking into consideration the background assumption – a more than minimal version of psychophysical parallelism – he makes, and which I share: for the impression of macroscopic change to emerge, Barbour, in accordance with our current best theories, relies on brain processes to happen. However, brain processes cannot be had from within individual, static configurations. While not logically being ruled out, time capsules alone don’t seem to provide an adequate solution to temporal solipsism. Accordingly, worlds should not be identified with Barbourian time capsules, but with sets of configurations that are sufficiently similar in the sense of not amounting to macroscopic differences for a given macroscopic object within a

configuration (IV.6). Before finally introducing worlds capable of forming FAPP-continuous world histories into the framework, I examine some preliminaries regarding the entities worlds are supposed to consist of, and the space they inhabit (IV.7). In particular, issues revolving around particle identity, impenetrability and particle number conservation will be discussed (IV.7.1). With the appropriate notion of worlds thus developed (IV.8), world histories, FAPP-continuously linking up remote worlds in configuration space and time, can emerge (IV.9), thereby solving the issue of temporal solipsism while being compatible with a more than minimal psychophysical parallelism (IV.10).

II. A short introduction to Minimal Bohmian Mechanics¹

To start off with a very quick overview, let me state the core features of Minimal Bohmian Mechanics. MBM can be considered a descendant of Bohmian Mechanics and the Everett interpretation (henceforth “Everett” for short) of non-relativistic Quantum Mechanics (QM). It is a Bohmian theory in the sense that it features a similar ontology as BM, namely particles endowed with definite locations in physical space at all times. These particles in 3-space are regarded as fundamental in a sense to be made explicit. Nevertheless, MBM does not necessarily fall under the primitive ontology approach, although it can be understood as a primitive ontology theory, as we will see shortly.² For convenience, I will often refer to the particle configuration represented by a point in 3N-dimensional configuration space, sometimes called the super-particle³, rather than to the N-particle system in 3-dimensional physical space. This does not imply some kind of ontological order between the two: The N-particle system in 3-dimensional physical space and the 1-

¹ Patrick Dürr (Oriol, Oxford) and I currently have a paper on MBM under review. For various reasons, I cannot fully reproduce it here. Instead, I will outline MBM in a significantly shortened and simplified manner appropriate for our present purposes. For a more detailed account, the reader is referred to the paper mentioned. Until publication, the manuscript is available upon request.

² The primitive ontology approach shall be explained in paragraph II.4. For reasons to be elaborated there, I am hesitant calling MBM’s ontology a primitive ontology *sensu stricto* without further qualification.

³ Or “universal particle” or “world-particle” in Albert’s terminology, cf. Albert (1996, 278).

super-particle system in $3N$ -dimensional configuration space are one and the same thing, merely represented differently. The super-particle, i.e. the configuration of the N -particle system as a whole, is actualized at a single point in configuration space. MBM deviates from orthodox BM in that neither particles in 3-space nor the super-particle move on continuous trajectories. Instead, configurations are actualized at random, following the quantum equilibrium hypothesis (QEH), understood in MBM as a physical law stochastically restricting the actualizations of the super-particle in accordance with the probabilities given by the Born rule (BR): The super-particle jumps through configuration space, following the probability density associated with the universal wave function (WF). The WF itself, as usually, evolves deterministically according to the standard Schrödinger equation (SE).

From these jumps of the super-particle through configuration space, together with the branching structure of the WF induced by decoherence, stems the main feature MBM inherits from Everett, namely its many-worlds character: at different times, the super-particle actualizes different configurations and, a fortiori, instantiates different world branches. The disposition for a given configuration to be actualized is encoded in the WF, with its probability – understood as a (single-case) propensity – given by the BR. MBM's actualizations are genuinely, irreducibly chancy, but weighted: individual actualizations are always probabilistically restricted by the QEH in such a way that they collectively approximate Born statistics.

A more detailed picture of MBM will emerge from a comparison with standard BM. Alterations of BM's formalism, and a reinterpretation of probabilities in terms of propensities, as just mentioned, lead to the formulation of MBM. MBM turns out to be empirically equivalent with orthodox QM, as required of every viable alternative theory. Before giving a quick outline of standard BM, however, a few remarks about terminology are in order:

It is important to be clear about which theory we refer to when speaking of (standard) Bohmian Mechanics. In particular, Bohmian Mechanics as put forward by

Dürr et al. (1992), despite its name, must not be confused with Bohm's original theory as set out in his 1952 papers (Bohm 1952a,b). The latter goes under different names. Sometimes it is called "the Bohm interpretation of quantum mechanics", or simply "the Bohm interpretation" (Hiley 2009, *passim*), sometimes "causal" or "ontological interpretation" (cf. Friebe et al. 2015, 195). Goldstein (2017, §5) refers to it as "the quantum potential formulation of the de Broglie-Bohm theory" or "the quantum potential formulation of Bohmian mechanics". Cushing heads chapter 4 of his 1994 book "Bohm's Quantum Theory", but usually refers to it simply as "Bohm's theory" (Cushing 1994, *passim*).⁴ I shall follow Cushing in calling this theory "Bohm's theory", or alternatively "Bohm's original theory", where the emphasis is called for. All these names refer to a theory that is different from what I, in accordance with Dürr et al. (1992), call Bohmian Mechanics, both formally and ontologically. Cushing (1994, §2.1, 9, *emphasis in original*) remarks that "a scientific theory can be seen as having two distinct components: its formalism and its interpretation", where "formalism means a set of equations and a set of calculational rules for making predictions that can be compared with experiment", and "interpretation refers to what the theory tells us about the underlying structure of these phenomena (i.e., the corresponding story about the furniture of the world – an ontology). Hence, *one* formalism with *two* different interpretations counts as *two* different theories". Adopting this distinction, Bohm's original theory and Bohmian Mechanics turn out to be decidedly different theories, for they differ in both formalism and interpretation. The former involves a quantum potential⁵ that contributes to a force acting on the particles (Bohm 1952a, 170): "the particles moves [sic] under the action of a force which is not entirely derivable from the classical potential [...] but which also obtains a contribution from the 'quantum-mechanical' potential". In Bohmian Mechanics, by contrast, the quantum potential plays no fundamental role. There are no fundamental forces acting on Bohmian particles (Goldstein 2017, §5): "Bohmian mechanics as presented here is a first-order theory, in which it is the

⁴ However, he also refers to it (*ibid.*, 47) as "the causal interpretation".

⁵ And restrictions on particle momenta to render the theory empirically equivalent with QM, see Bohm (1952a, 170), Cushing (1994, 43 and 45).

velocity, the rate of change of position, that is fundamental. It is this quantity, given by the guiding equation, that the theory specifies directly and simply. The second-order (Newtonian) concepts of acceleration and force, work and energy do not play any fundamental role. Bohm, however, did not regard his theory in this way. He regarded it, fundamentally, as a second-order theory, describing particles moving under the influence of forces, among which, however, is a force stemming from a ‘quantum potential’.” Considering their different interpretations, i.e. their different ontologies, it becomes obvious that Bohm’s original theory and Bohmian Mechanics are not the same theory. The former, with its quantum potential generating a force on particles, provides “a story about the furniture of the world” quite different from the story told by the latter. As reported by Hiley (1999, 117), Bohm agrees with this assessment: “It should be noted that the views expressed in our book [Bohm & Hiley 1993] differ very substantially from those of Dürr et al. (1992) who have developed an alternative theory. It was very unfortunately that they chose the term ‘Bohmian mechanics’ to describe their work. When Bohm first saw the term he remarked, “Why do they call it ‘Bohmian mechanics’? Have they not understood a thing that I have written?” [...] It would have been far better if Dürr et al. (1992) had chosen the term ‘Bell mechanics’. That would have reflected the actual situation far more accurately.” Indeed, a lot of terminological confusion could have been avoided. However, that horse has left the barn. Goldstein et al.’s theory is standardly referred to as Bohmian Mechanics, and I will adhere to that terminology.

II.1 Bohmian Mechanics as a solution to the measurement problem

BM sets out to complete the formalism of QM and thus to provide a solution to what is generally referred to as “the measurement problem” of QM, arising from the general presence of superposition states.⁶ Maudlin (1995) provides a concise statement of, as he calls it (*ibid.*, 7), the “problem of outcomes”. He presents three

⁶ This is not to say that the solution of the measurement problem was the motivation behind the development of BM from the historical beginning, but merely that it is one of the main motivations for adopting BM today.

propositions at most two of which can be true at the same time. The fact that the conjunction of these propositions is inconsistent constitutes the measurement problem. Paraphrasing Maudlin:

- (A) The WF provides a complete description of the physical properties of a system
- (B) The WF always evolves according to the SE
- (C) Measurement outcomes are always determinate.

To see why (A), (B) and (C) can't all be true, think of the following example of the measurement of the spin of an electron, in a certain direction. The outcome of such a measurement is determinate if (ibid.) "at the end of the measurement the measuring device is either in a state which indicates spin up (and not spin down) or spin down (and not spin up)" in that direction.

Before the measurement, the device is in its pre-measurement ready state, $|ready\rangle_d$.⁷ Let's say it is a device that measures spin in z-direction, and we prepare an electron with spin up in that direction, written as $|\uparrow_z\rangle_e$. The expression $|\uparrow_z\rangle_e \otimes |ready\rangle_d$ denotes the state of the total system, consisting of the electron and the measuring device, before the measurement. After the measurement, the state of the total system will be $|\uparrow_z\rangle_e \otimes |up\rangle_d$, with $|up\rangle_d$ indicating that the measuring device shows the outcome "up", as expected of a properly working device.

If we prepare the electron to be in the state $|\uparrow_x\rangle_e$ (spin up in x-direction), say, it will be in a superposition of z-spin states, which can be expressed as $|\uparrow_x\rangle_e = 1/\sqrt{2} |\uparrow_z\rangle_e + 1/\sqrt{2} |\downarrow_z\rangle_e$. Accordingly, the pre-measurement state of the total system is

$$|\uparrow_x\rangle_e \otimes |ready\rangle_d = (1/\sqrt{2} |\uparrow_z\rangle_e + 1/\sqrt{2} |\downarrow_z\rangle_e) \otimes |ready\rangle_d$$

Proposition (B) above tells us that this pre-measurement state of the total system evolves into the post-measurement state $1/\sqrt{2} |\uparrow_z\rangle_e \otimes |up\rangle_d + 1/\sqrt{2} |\downarrow_z\rangle_e \otimes$

⁷ Except for minor adaptations, I follow Maudlin's example and notation.

$|down\rangle_a$. But being in this state, the measuring device isn't in a determinate state of either showing the outcome "up" or the outcome "down". It's in a superposition. According to (A), this is a complete description of the physical properties of the system. So, assuming that (A) and (B) are true, the evolution of the system according to the SE results in the system being in an indeterminate state, and this state is the complete description. This contradicts proposition (C).

Ways out of this predicament are to abandon one of the three propositions. Everettians will deny that (C) is true. Proponents of collapse⁸ theories reject (B). Bohmians hold that the WF isn't a complete description of the physical properties of a system. They thus abandon (A) and attempt to complete the description with a definite particle configuration.

BM, with its completion of the state description by adding a definite particle configuration, indeed can provide a solution to the measurement problem. In light of the severity of this problem, this provides good motivation to consider Bohmian theories as viable alternatives to orthodox QM and its various interpretations. However, for a theory, or a class of theories, to be viable alternatives to a given theory, it is not sufficient to solve some problem of the latter, be that as important as it may. Being a viable alternative requires empirical equivalence with the theory to be replaced. In the following sections, I will therefore first provide a general description of BM and establish its empirical equivalence with orthodox QM, and then give an account of MBM's empirical equivalence with BM. Via the transitivity of empirical equivalence, MBM can be said to be empirically equivalent with QM, iff it is empirically equivalent with BM, and BM is empirically equivalent with QM. So let's first describe BM and its empirical equivalence with QM in more general terms.

⁸ As Maudlin notes (*ibid*, 8), it would be better to speak of non-linear evolution instead of collapse, since there are theories in which the evolution of the WF is always non-linear. However, so-called collapse interpretations are the most prominent amongst those featuring a non-linear evolution of the WF.

II.2 A general description of BM and its empirical equivalence with QM

BM is a non-relativistic quantum theory. It describes a universe consisting of N particles and their deterministic, continuous evolution in space. The particle configuration Q_t at time t is given by the definite particle positions of the N particles at that time. BM adheres to the standard, non-relativistic SE, describing the evolution of the universal quantum wave function Ψ for a universe of N particles. The guiding equation (GE), featuring Ψ , describes the evolution of the universal particle configuration Q in terms of continuous trajectories of the Bohmian particles. Given some initial configuration Q_{t_0} , i.e. the initial spatial positions of all N particles at initial time t_0 , the GE determines the configuration for all times t_i , before and after t_0 . The quantum equilibrium hypothesis (QEH) states that (initial) ensembles of identically prepared subsystems of the universe, each with the WF ψ , typically, i.e. almost always, are in fact Born distributed with theoretical distribution $\rho_{th} = |\psi|^2$, approximating the empirical distribution ρ_{emp} for large ensembles of identically prepared subsystems.⁹ This requires that ours is a typical universe, and that in typical universes, typical ensembles are Born distributed.

Justifying the typicality approach to BM is a rather intricate matter. I am unable to engage with it here. However, the alleged typicality of our universe notwithstanding, since BM is a deterministic theory and contains the assumption that particle positions are always definite, its probabilities for finding an individual particle at a definite location are often construed as a measure of the ignorance of an observer about the precise locations of particles, especially their initial distribution.¹⁰ While the typicality approach explains why one finds the distribution

⁹ Cf. Goldstein (2012, 66): “The theoretical distribution is an idealization providing a good approximation to the empirical distribution, $\rho_{emp} \approx \rho_{th}$, in the limit of large ensembles of subsystems.”

¹⁰ See, e.g., Bell (1987b, 112): “the only use of probability here is, as in classical statistical mechanics, to take account of uncertainty in initial conditions.” Likewise, Passon (2010, 53): “The result of the experiment is random only for us, because we don’t know the initial conditions.” (My translation.) Callender (2007, 364) acknowledges this fact – “When speaking of probabilities in a deterministic

of particles in *ensembles* of identically prepared systems as $|\psi(Q)|^2$, it does not do away with the ignorance interpretation of probabilities for the position of particles in *individual* systems. For them, an observer can only provide probabilities to be here or there, because they don't *know* where the particles *actually* are. But since, according to BM, they must be at a definite location, which evolves deterministically according to the GE, these probabilities do not reflect a fundamentally stochastic nature, but the observer's ignorance.¹¹

Especially if one regards BM to be a "quantum theory without observers" (cf. Allori et al. 2008, *passim*), it *prima facie* seems unclear whose ignorance about the initial configuration it is that is reflected in BM's QEH. Also, shouldn't probabilities be construed as fundamental and objective in a theory that features them at its heart and claims to be objective, without observers? My intuition tells me that they should. But irrespective of whether your intuition agrees with mine on that matter, it is clear that on MBM's irreducibly stochastic actualization dynamics, probabilities are construed as fundamental and objective. If you share my intuition, you might regard this as an advantage MBM has over BM.

QEH and the GE together guarantee empirical equivalence between BM and QM. For example, the particle positions measured at the screen of a repeated double slit experiment will (almost certainly) approximate the distribution expected according to BR. The GE ensures that this is always the case: *If* an ensemble of initial

theory one often hears it said that the probabilities are 'epistemic,' i.e., due to ignorance" – but rejects the ignorance interpretation.

¹¹ With respect to this ignorance interpretation of probabilities, BM resembles Bohm's original theory, (cf. Bohm 1952a, 166): "In contrast to the usual interpretation, this alternative interpretation permits us to conceive of each individual system as being in a precisely definable state, whose changes with time are determined by definite laws, analogous to (but not identical with) the classical equations of motion. Quantum-mechanical probabilities are regarded (like their counterparts in classical statistical mechanics) as only a practical necessity and not as a manifestation of an inherent lack of complete determination in the properties of matter at the quantum level." As the reason for this practical necessity, he cites "uncontrollable disturbances" (*ibid.*, 171): "In the usual interpretation, however, the need for a probability description is regarded as inherent in the very structure of matter [...], whereas in our interpretation, it arises [...] because from one measurement to the next, we cannot in practice predict or control the precise location of a particle, as a result of corresponding unpredictable and uncontrollable disturbances introduced by the measuring apparatus."

configurations already approximates the Born distribution, it will also approximate the Born distribution at earlier and later times, because the dynamics given by the GE preserves the initial distribution. Hence, in the experiment just mentioned, if the ensemble was (approximately) Born distributed at its preparation, it will be at the time of measurement at the screen, and vice versa. In other words, regarding observable phenomena, BM comprises the same claims as QM does: the theories are empirically equivalent.

Apart from their positions, the usual dynamical properties of particles – the properties that can change over time – are contextual in BM. Spin is an example of such a contextual property of a particle, a quantity that supervenes on the complete state description WF + configuration. It is a way of describing the particle's behaviour in certain contexts. Accordingly, the spin of a particle is determined by its behaviour under certain experimental conditions, e.g. in a Stern-Gerlach experiment. Recall that also in orthodox QM, a spin measurement ultimately is a position measurement: one of two detectors registers an event, while the other doesn't. Whether this is interpreted as a Bohmian particle ending up in said detector, or as a wave function collapsing onto the corresponding eigenstate upon measurement, makes a difference in metaphysics and ontology, but, empirically, amounts to the same phenomenon. Again, the theories exhibit empirical equivalence.

So much for this quick outline of BM and its empirical equivalence with QM. It shall provide the background against which, in the following section, MBM is discussed. In particular, I will establish MBM's empirical equivalence with BM, and a fortiori with QM.

II.3 From Bohmian Mechanics to Minimal Bohmian Mechanics

In essence, MBM is the somewhat elaborate claim that the GE is not strictly necessary for a Bohmian theory, i.e., that a Bohmian theory without a GE is viable,

in the sense that it is internally consistent and empirically equivalent to standard BM – and hence orthodox QM – and metaphysically adequate.

To be more precise, MBM, formally, retains the SE (including the WF) and QEH from BM, but not the GE. In addition, it reinterprets QEH as a stochastic law for the actualization of configurations. Together, the SE, the WF and QEH guarantee empirical equivalence: the SE provides the dynamics for the WF, and the WF as featuring in the QEH restricts the distribution of particles in ensembles, in the same way it did in BM for the so-called “initial” distribution. But while in BM, the role of the QEH is most obvious for this initial distribution, MBM makes its role explicit for all distributions, at all times. Recall that, in BM, particle positions are evolved by the GE in such a manner that the distribution is preserved. If an ensemble is Born distributed once, it is for all times.¹² But note that, under the condition that QEH is assumed for all times instead of merely some initial time t_0 , the WF, SE and QEH alone, i.e. without the GE, provide sufficient theoretical structure to guarantee Born distribution at all times. More colloquially: If we feel justified to assume QEH for some *arbitrary* point in time, why shouldn’t we feel justified to assume it for *any* point in time, calling the distribution at this time the “initial” distribution, if so desired? Of any distribution, QEH will tell us that it approximates Born statistics. One might be compelled to reply that it is simpler to assume QEH just for one instant and let the dynamics do the work. But note that this requires an additional assumption itself, namely the deterministic particle dynamics as represented by the GE. In effect, we are replacing two assumptions – that QEH holds for the initial configuration, and that the GE evolves the configuration from there on, such that

¹² Not taking into account the possible relaxation of non-equilibrium probability distributions into equilibrium as shown by Valentini, which requires the assumption that the initial probability distribution is suitably coarse-grained (Valentini 1991, 8): “The proof rests on the assumption of no ‘micro-structure’ for the initial state, as for the classical case. Specifically, we assume the equality of coarse-grained and fine-grained quantities at the initial time $t=0$, i.e. we assume $\bar{P}(0) = P(0)$, $|\bar{\psi}(0)|^2 = |\psi(0)|^2$.” As Norsen (2018) shows, this allows for the evolution of the probability distribution not only from non-equilibrium to equilibrium, but also from equilibrium to non-equilibrium. Given we do not assume QEH and instead follow Valentini’s derivation, it cannot be claimed that an ensemble is Born distributed at all times, if it is Born distributed once. However, for MBM, we assume QEH, such that the claim still holds.

the ensemble always approximates Born distribution – by one assumption, namely that the QEH holds for all configurations, at all times. Arguably, the exchange is worthwhile: While avoiding the assumption of a suitable GE, including the metaphysical baggage this brings with it, we merely extend the assumption that the QEH holds for (at least) one instant to the assumption that it holds for all instants. As long as there is no strong reason against it, this seems fair. After all, standard BM includes the same claim of the ensemble always being (approximately) Born distributed, if it ever is. The difference between BM and MBM is how we arrive there: MBM assumes applicability of the QEH at all times. BM arrives there via the assumption of the applicability of the QEH at some (initial) time *plus* the assumption of an appropriate particle dynamics (GE) that guarantees Born distribution at all other times as well, given Born distribution at the initial time.

Let's elaborate. Does the GE play an indispensable role in the preservation of Born distribution over time? According to MBM, it doesn't. In BM, the GE serves a dual purpose: The first is to evolve the positions of individual particles *deterministically*. The second is to do this in such a manner that the ensemble distribution at all times is compatible with the restrictions imposed by QEH, given the initial distribution is. QEH in BM is exactly that, "the assumption that the initial ($t = 0$) configuration of a quantum system was random with distribution P_B [i.e. Born distribution]." (Norsen 2018, 5) Regarding its first purpose, it is easy to dispense with it, if one doesn't insist on determinism.¹³ Regarding its second purpose, within the framework of MBM, the GE is unnecessary as well. Taking the QEH as restricting not only the *initial* configuration, but the random actualizations of configurations in general, ensembles of actualized configurations will (approximately) show the required Born distribution. Postulating trajectories, then, is superfluous. Following Norsen's (2018)

¹³ Even proponents of BM, despite regarding BM as (Goldstein et al. 2011, 12) "the simplest satisfactory version of quantum mechanics", don't feel particularly committed to determinism: "The circumstance that it uses point particles and is deterministic does not mean that we are dogmatically committed to point particles or determinism; also indeterministic theories with another ontology may very well be satisfactory." In the same spirit, they (Dürr & Teufel 2009, 9) reject the statement that "the aim of Bohmian mechanics is to restore determinism in the quantum world. That is false. [...] What is 'out there' could just as well be governed by stochastic laws [...]. It happens to be deterministic, which is fine, but not an ontological necessity."

presentation of equivariance, one could say that with the equivariance of $\rho = |\Psi|^2$, “the usual quantum ‘probability density’” (ibid., 4) or “wave function ‘intensity’” (ibid., 5), the equivariance of the “particle probability distribution” P_B in MBM is ensured, quite trivially, via the QEH. Additional dynamics preserving the initial distribution over time isn’t necessary. Either way, equivariance is achieved, and it is equivariance that is responsible for the fact that the actual, empirical particle distribution “tracks” the probability density. Because both theories have this feature, they are empirically equivalent, and a fortiori, because of the empirical equivalence of BM and QM, MBM and QM are empirically equivalent as well. In conclusion, for empirical equivalence with QM, trajectories are not needed.

This argument by itself might not convince everyone to adopt MBM over standard BM. One might have extra-empirical reasons to prefer a theory with continuous, deterministic trajectories, despite them not being strictly necessary. In particular, one might argue with Bell that a fundamentally stochastic dynamics of the configuration leads to a temporal form of solipsism. This issue will be discussed at length. For now, I am not going to take sides on the matter. My present claim is merely that MBM is a viable alternative to BM, given one doesn’t insist on a deterministic, continuous particle dynamics, because empirically, the two theories are equivalent.

How, then, do particle configurations in MBM evolve in such a manner that QEH holds at all times? This is due to QEH restricting the random, discontinuous particle dynamics that replaces the GE. In MBM, the particle configuration – i.e. the super-particle – of the universe is assumed to instantaneously jump through $3N$ -dimensional configuration space, in accordance with QEH. Accordingly, the N Bohmian particles locate and dislocate themselves in 3-dimensional physical space. Without requiring any external triggering mechanism – after all, the universe by definition is all there is, such that there is nothing external to it – these jumps of the super-particle are spontaneous. The disposition for certain configurations to be actualized is encoded in the WF, and the respective probabilities (propensities) as

given by BR – being a measure of a fundamental feature of the universe, namely said disposition – are objective. In contradistinction to BM’s deterministic particle dynamics, actualizations in MBM at any time are fundamentally random. They are stochastically independent from each other: $P(Q_t \cap Q_{t'}) = P(Q_t) * P(Q_{t'})$.¹⁴ There is no link between actualized configurations at different times. In Bell’s (1987b, 135) words: “there is no association of the particular present with any particular past.” But as has been argued, this link, which would be provided by trajectories in BM, is not necessary for empirical equivalence with orthodox QM. The stochastic dynamics of the super-particle, following the probability density as it jumps, is sufficient, because it adheres to the restrictions imposed by QEH, and being in accord with QEH at all times is sufficient for being empirically equivalent with QM.

Having established the empirical equivalence of MBM, BM and QM, despite the former’s lack of continuous trajectories, I shall now turn to two interrelated issues: MBM’s classification as a primitive ontology theory, and its solution to the problem of communication, arising from the two-space reading of BM. Many proponents of BM deem a theory’s classification as a PO crucial for it being bona fide Bohmian. In light of this, we must ask whether MBM can be made compatible with the PO

¹⁴ Some remarks are in order regarding the stochastic independence of MBM’s actualizations at different times. At first blush, the stochastic independence of actualizations may seem counterintuitive. Underlying this impression is the following thought: the probability density distribution over configuration space is determined by the WF. The latter evolves deterministically, governed by the SE, so the probability density at one time isn’t independent of the probability density at another time. The actualization at one time might be informative about the probability density at that time, and via the latter’s deterministic evolution, about the probability density at other times. Thus, one might be tempted to conclude, the actualization of $Q_{t'}$ at t' isn’t stochastically independent of the actualization of Q_t at t . This reasoning is misleading, however. In responding as follows, I hope to make clear that MBM’s actualizations indeed are stochastically independent.

On MBM, there is no collapse of the universal WF upon the actualization of a configuration. If there was, an actualization would suddenly change the WF and hence the probability density distribution for subsequent actualizations. Because there is no such collapse, the probability density distribution governing the actualization of $Q_{t'}$ at t' isn’t influenced by the actualization of Q_t at t . We may therefore take the probability of the joint event to factorise. The stochastic independence of actualizations is not affected by the lawlike dependence between probability density distributions at different times. The latter are related to each other via the deterministic evolution of the WF, governed by the SE. But this has no bearing on the fact that actualizations at different times don’t affect each other.

I thank Jiji Zhang (Lingnan) for pressing me on clarifying this point.

approach. After presenting the PO approach, I will argue that this is indeed possible, but that, in order to also solve the problem of communication, our interpretational freedom with respect to the WF is restricted. In particular, subsuming MBM under the PO paradigm seems to require abandoning the view that the WF encodes a disposition.

II.4 MBM as a PO theory and the problem of communication

As just noted, MBM *can* be understood as a primitive ontology theory, when imposing certain restrictions on the interpretation of the WF. In principle, at least in my view, nothing much hinges on this classification. However, since proponents of BM often regard applicability of the PO framework as crucial for a theory to be Bohmian, I want to show that MBM can be made compatible with it, thereby preempting the criticism that MBM is not a viable alternative to BM because it cannot be classified as a PO theory. Abandoning an ontologically thick, dispositionalist view of the WF (and instead adopting an ontologically thin stance towards the WF) opens up this possibility.

II.4.1 What makes a PO theory?

To tackle this issue, we first need to be clear about the central features of the PO approach. What makes a PO theory? Following the literature, the PO approach imposes a strict distinction between primitive and non-primitive elements of the ontology of a theory. It is only the primitive elements that are fundamental and make up/compose non-fundamental physical objects. Allori (2015, 107f.) writes: “what constitutes physical objects, the PO, [...] lives in three-dimensional space or space-time and constitutes the building blocks of everything else. In the formalism of the theory, the variables representing the PO are called the primitive variables. In addition, there are other [non-primitive] variables necessary to implement the dynamics for the primitive variables”. Pertaining to fundamental theories, the

primitive variables (ibid., 109) “are ontologically primitive, given that they represent matter and they provide the fundamental entities the theory describes.”

In order to make the notion of PO more precise, this fundamentality claim should be spelled out. It consists of two separate claims: (1) The PO is fundamental in the sense that all other physical objects are composed of its elements. (2) The PO is fundamental in the sense that it itself is not ontologically dependent on further, non-3D entities.

I shall refer to the kind of fundamentality claimed in (1) as $\text{fundamental}_{\text{CMB}}$: The primitive entities are $\text{fundamental}_{\text{CMB}}$ in that they form the *complete minimal basis* (CMB) for composed physical objects.¹⁵ That is to say, the PO contains all and only those 3D entities from which all other physical objects are made. The qualifications “complete” and “minimal” are of some importance. The former (“complete”) specifies the PO such that it contains *all* the elements necessary to make up other physical objects. The reason for imposing this completeness requirement should be immediately obvious: if the PO were not complete, not all other physical objects could be composed from its elements. The latter (“minimal”) specifies the PO such that it contains *only those* elements. I.e., the PO does not contain entities beyond those necessary to compose all other physical objects. The reason for imposing this requirement is fairly obvious, too: Without insisting on the PO being minimal in order to be fundamental, we could simply regard the set of all physical objects, including composed objects, to be the PO.¹⁶ With the PO being fundamental, this would render composed objects fundamental as well. At best, this amounts to an empty notion of fundamentality. At worst, it is an outright contradiction. In any case, it is at odds with the PO approach.

¹⁵ I thank an anonymous reviewer of my paper with Patrick Dürr for pressing on this further specification of the PO in terms of a complete minimal basis.

¹⁶ Cf. Tahko (2018, 1.3): “We have included a *minimality* condition in the definition of (CMB). This is an important addition, because otherwise we could take the set of all the entities in the world and call it complete, given that this set as well would include all the fundamental entities. So, according to (CMB), the complete minimal basis must include all and only the fundamental entities.”

The kind of fundamentality claimed in (2) shall be referred to as $\text{fundamental}_{\text{ODEP}}$. Since it contains the notion of ontological dependence on non-3D entities, we must flesh out the meaning of this notion. Esfeld & Oldofredi (2018, 11) demand that the elements of the PO be fundamental (i.e., $\text{fundamental}_{\text{ODEP}}$) in the sense that “they are not reducible to more elementary notions”. Ontological independence, then, simply means to be irreducible in this sense. However, this seems to be not precise enough, for ontological independence in terms of irreducibility can mean several things. For example, an entity can be said to be irreducible if it isn’t instantiated by a non-3D entity playing its functional role¹⁷, or if it isn’t identical to a non-3D entity.¹⁸

Demanding irreducibility only with respect to *non-3D* entities requires an explanation. Arguably, it already would be a blow to a PO theory if its PO should turn out to be reducible to *3D* entities, and hence to “more elementary notions”. But that blow wouldn’t be fatal. It would still be possible to revise the theory, such that it contains as its new PO those 3D entities to which its former PO turned out to be reducible. As long as the new PO is not reducible any further, it can be regarded as $\text{fundamental}_{\text{ODEP}}$. The revised theory, then, remains a PO theory. Reducibility to 3D entities thus might be an inconvenience, but it doesn’t hinder the PO approach in principle. However, reducibility of the PO to non-3D entities is fatal to the PO approach. In this scenario, the theory can’t simply be revised such that it contains as its new PO those entities that turned out to be fundamental, because the PO approach doesn’t allow for non-3D entities in the PO. Hence, demanding

¹⁷ On Everett, particles, just like macroscopic objects, are reducible in this sense: They are structural patterns, instantiated by the (non-3D) universal WF, playing the functional role of such objects. As such, they are useful ways of describing reality, albeit at a non-fundamental level (cf. Wallace 2012, Ch. 2). Fundamentally, all there is is the universal WF.

¹⁸ In addition to these options, one might be tempted to spell out reducibility in terms of supervenience. An entity, then, would be reducible if it supervenes on a non-3D entity, i.e., if change in the former cannot be had without change in the latter. However, while supervenience seems necessary for reducibility, it is unclear whether supervenience suffices for reducibility, such that the latter could be spelled out in terms of the former. As McLaughlin & Bennett (2018, 3.3) put it: “Everyone agrees that reduction requires supervenience. [...] The more interesting issue is whether supervenience suffices for reduction”. It lies outside the scope of this thesis to engage with this debate. I will therefore not further discuss supervenience as a candidate sufficient for spelling out reducibility of the PO.

irreducibility only with respect to non-3D entities seems apt: the PO must be ontologically independent in that it is not reducible to non-3D entities, e.g. the 3N-dimensional WF. That is, no non-3D entity may be playing the PO's functional role or be identical with it.

In light of this restricted demand, there is another type of ontological dependence that should be mentioned, because it plays a somewhat special role in the context of classifying theories as PO theories. One could argue that entities are not $\text{fundamental}_{\text{ODEP}}$ if they have proper mereological parts. That is to say, $\text{fundamental}_{\text{ODEP}}$ entities should be mereologically simple. At first blush, a PO not being mereologically simple seems to be fixable just as easily as general dependence on 3D entities: if elements of the PO turn out to have proper mereological parts, the theory can be revised, such that the new PO consists of those proper mereological parts. The PO still wouldn't be reducible to non-3D entities. However, an issue arises if there can't be mereological simples in the PO, such that it can never be fundamental. Consider the GRW matter density theory (GRWm), usually classified as a PO theory: with the matter density being continuously spread over physical space, every mereological part of it has further, proper mereological parts. As Esfeld (2018, 174) diagnoses, "the matter density stuff, being gunk, is infinitely divisible". Being infinitely divisible, the matter density cannot be $\text{fundamental}_{\text{ODEP}}$, not because it itself has proper mereological parts, which would be fixable by revising the theory such that the parts are regarded as the PO, but because its proper mereological parts have proper mereological parts as well, and so on, ad infinitum. Thus, it is impossible to arrive at a revision featuring a truly $\text{fundamental}_{\text{ODEP}}$ PO. To resolve this issue, we can either decide to not regard GRWm as a PO theory, or to allow for $\text{fundamental}_{\text{ODEP}}$ entities not being mereologically simple. I leave this for the proponents of the PO approach to decide. Since MBM's particles aren't infinitely divisible, this is not an issue for the present investigation.

To sum up: the primitive ontology of a theory consists of those and only those 3D entities that are $\text{fundamental}_{\text{CMB}}$ – i.e. that form a complete minimal basis – to all other physical objects, and that are $\text{fundamental}_{\text{ODEP}}$, i.e. that are not reducible to other, non-3D entities, in the sense elaborated.

II.4.2 Is MBM compatible with the PO approach?

Prima facie, MBM's particles aren't reducible to non-3D entities: they are not instantiated by non-3D entities playing their functional role, aren't identical to non-3D entities, and don't have proper mereological parts. That is to say, they are $\text{fundamental}_{\text{ODEP}}$. In this respect, MBM seems compatible with the PO approach. At least for now.

But are MBM's particles $\text{fundamental}_{\text{CMB}}$ as well? This depends on the ontological status and role ascribed to MBM's WF. If the latter, under a given interpretation, turns out to be part of the complete minimal basis and hence of the primitive ontology, then, under this interpretation, MBM cannot be a PO theory, since the PO approach doesn't allow for non-3D entities in the PO, which the WF is.

An interpretation of the WF, however, must not simply be chosen at will, such that the theory can be classified as a PO theory. A further requirement is imposed, namely that the theory be able to solve the “problem of communication”¹⁹ arising from the (Ney 2012, 11) “two-space reading of Bohmian mechanics” (and likewise of MBM), i.e. arising from the fact that (ibid.) “the theory posits the existence of two physical spaces: a high-dimensional configuration space occupied by the wavefunction, and a separate, three-dimensional physical space occupied by the many particles that make up our tables, chairs, and the other objects of our manifest image of the world.” In short, to render MBM a PO theory, the interpretation of its WF must meet two demands at the same time: it must solve or

¹⁹ As it is called e.g. in Suárez (2015) and references therein.

circumvent the problem of communication while not considering the WF to be part of the complete minimal basis. This, I hold, can only be achieved under a “nomological” interpretation, i.e. when the WF is regarded as a law, or lawlike. To see this, let’s have a look at the problem of communication.

II.4.3 The problem of communication

In standard BM, the N individual Bohmian particles, with position being their only non-contextual dynamical property, are located in ordinary, 3-dimensional physical space, and make up all physical objects, while the WF lives on $3N$ -dimensional configuration space, restricts their initial distribution and guides the evolution of their positions via the GE. This is the two-space reading of BM – particles living in 3-dimensional physical space, the WF living in $3N$ -dimensional configuration space – ultimately resulting in the problem of communication:

On the one hand, the WF’s role in the theory suggests interpreting it as a real, physical object, e.g. a field. After all, how can an abstract, purely mathematical object, defined on likewise abstract and purely mathematical configuration space influence fundamental, real objects, i.e. particles, in physical 3-space? To explain this influence, the WF should be considered a real part of the complete, physical state description and its evolution, at least in the sense that it has an effect on the fundamental stuff (Ney 2012, 11): “For even on this two-space reading of Bohmian mechanics, where the particles are ontologically fundamental, it is still true that it is states of the wavefunction that determine these particles’ behavior over time. The particles themselves only have positions. And what they do at later times according to this theory, i.e. what positions they move to over time, depends on the wavefunction.” Featuring in the GE, the WF guides the movement of particles on their trajectories through physical space (in a non-local manner). Likewise, it restricts their (initial) distribution, as it features in QEH. The central role the WF plays in the formalism, and for the empirical adequacy of BM, renders it difficult to immediately dismiss it as a purely mathematical artefact without any ontological

significance. As Bell (1987a, 128, his emphasis) writes: “*No one can understand this theory until he is willing to think of ψ as a real, objective field rather than just a ‘probability amplitude’. Even though it propagates not in 3-space but in 3N-space.*”²⁰

On the other hand, how to account for the influence this real physical object has on the fundamental particles, given the different spaces they live in? How does a physical object in 3N-dimensional configuration space tell the stuff in 3-space what to do? This is the problem of communication. As Suárez (2015, 10f.) notes, “this is generally a problem for any account on which the wave-function in configuration space ‘dictates’ features of physical 3D space that are responsible for the motion of particles, whether it be a multifold, forces acting in the space, a quantum potential located in that space, or directly the motions of the 3D particles themselves. [...] How does the wavefunction in configuration space *guide* the particles in physical space?”

There are several ways MBM can deal with this problem of communication. For example, one can regard the super-particle in configuration space as the fundamental_{ODEP} referent of MBM’s ontology while not claiming such fundamentality for the Bohmian particles in 3-space. Such a move would immediately solve the problem (by circumventing it): adopting Albert’s super-particle view, the fundamental_{ODEP} referent of the theory is put into the same space as the WF, i.e. 3N-dimensional configuration space. In contradistinction to the “two-space reading of Bohmian mechanics”, this can be called the one-space reading of MBM. The super-particle actualizes instantaneous, definite configurations in configuration space, with the actualization disposition encoded in the WF, living on the very same space. The problem of an entity living in one space influencing another, living in a different space, no longer arises. At the same time, through the configuration being definite at all times, particle positions in 3-space are definite as

²⁰ Note that this remark applies to BM as well as to Bohm’s original theory with a quantum potential (cf. Suárez 2015, 10f., as quoted in the following passage.) Bell himself, in his 1987a as well as in his unpublished paper from 1971 (see II.5 below), refers to Bohm’s work of 1952.

well, allowing for the composition of physical objects. Bohmian particles in 3-space, hence, remain fundamental_{CMB}. However, if we want MBM to be compatible with the PO approach as presented above, this option leads to a dead end: while, with adopting the super-particle as the fundamental_{ODEP} referent of the theory, Bohmian particles in 3-space remain fundamental_{CMB}, they are no longer fundamental_{ODEP}, for they either are secondary to the super-particle, or, at best, identical with it. Also, the super-particle being fundamental_{ODEP} conflicts with the PO approach, demanding that the PO should be entirely 3-dimensional. This can only be resolved if we think of the super-particle not as the fundamental referent of the theory, but as a mere representational variation, describing the configuration of the Bohmian particles in physical 3-space, such that the latter remain fundamental_{ODEP} and fundamental_{CMB}. But then, the solution of the problem of communication via the super-particle view is blocked, for we arrive at two distinct spaces after all.

A possible solution to *this* problem, of course, would be to relax the requirements imposed on PO theories, such that either non-3D entities may figure as primitives, or that 3D entities are not required to be fundamental_{ODEP}. However, construed as above, this is incompatible with the PO approach.

II.4.4 Disposition or law?

So, taking the super-particle to be MBM's fundamental referent is of no avail to the problem of communication. We hence shall stick with the usual Bohmian particles. In this case, the second option for solving the problem of communication must be sought in the interpretation of the WF. In what follows, I will consider dispositional as well as nomological interpretations. So far, dispositionalism was assumed for MBM's actualizations of particle configurations, more or less implicitly, without pointing to an alternative reading. To make up for this, let me be explicit now: dispositionalism_{MBM} is the view that, at any given time, the holistic N-particle system has a certain disposition, encoded in the WF, to actualize certain configurations. Note that the N-particle system is holistic in virtue of the particles

sharing this disposition: they aren't completely independent of each other. Rather, the N-particle system exhibits a genuinely quantum-mechanical feature – entanglement – that doesn't supervene on the individual properties and spatial distribution of the individual particles. In other words: the N-particle system as a whole has a feature that can't be reduced to the features of the individual particles. The latter, thus, aren't ontologically independent of the former. On dispositionalism_{MBM}, the problem of communication is straightforwardly avoided: the 3N-dimensional WF encodes the disposition for the actualization behaviour of the N-particle system as a whole, represented as a point in 3N-dimensional configuration space. But because of that, dispositionalism_{MBM} isn't compatible with the PO approach: the particles in 3-space aren't fundamental_{ODEP}, and the super-particle isn't 3-dimensional. Neither are suitable candidates for a primitive ontology.

When adopting dispositionalism_{MBM}, it seems, MBM isn't compatible with the PO approach. But dispositionalism certainly isn't the only option to interpret the WF. Chen (2019) provides a survey of several realist interpretations of the WF, amongst them nomological interpretations (ibid., 8): “nomological interpretations hold that the wave function is nomological, i.e., on a par with laws of nature.” He distinguishes “two kinds of nomological interpretations of the wave function: the strong nomological interpretations and the weak nomological interpretations.” Dürr et al. (2012, 12.3) discuss the possibility of the former for BM (ibid., 266): “one should think about [...] the possibility that it's nomological, nomic – that it's really more in the nature of a law than a concrete physical reality.” Likening it to equations of motion in classical mechanics, they suggest an analogous interpretation of the WF (ibid., 267), as “just a convenient device in terms of which the equations of motion can be nicely expressed. We're suggesting that you should regard the wave function in exactly the same way.” Esfeld's “quantum Humeanism”, following Chen, is an example of a weak nomological interpretation. Esfeld (2014, 455) holds that it is the primitive ontology approach to quantum physics that first of all enables quantum Humeanism: “the background that enables

Humeanism to stand firm is the development of what is known as primitive ontology theories of quantum physics.” Apart from the matter distributed in ordinary space(-time), there is, according to Esfeld, nothing but a law “as that what fixes (in a probabilistic or a deterministic manner) the temporal development of the distribution of matter in physical space, given an initial configuration of matter. That’s all. In particular, the quantum mechanical wavefunction is part and parcel of the law instead of being a physical entity on a par with the primitive ontology.” For his Humean approach, however, this law must not exist independently of the PO, but supervene on it (ibid., 456): “the Humean has to regard the law as supervening on the distribution of matter throughout the whole of space-time, that is, the entire mosaic of ‘local beables’ or local matters of particular fact.” As such, it is part of the “best system” systematizing these “local matters of particular fact” (ibid., 459): “For the Humean, the universal wavefunction and the dynamical law that appear in a physical theory such as Bohmian mechanics or the GRW theory are part of the best system, that is, the system that achieves the best balance between being simple and being informative in capturing what there is in the physical world. That system – and everything that belongs to it – supervenes on the entire distribution of the local matters of particular fact throughout the whole of space-time.”

It is neither my intent to arbitrate between different options of nomological interpretations of the WF, nor to advocate nomological interpretations in general. That being said, for those insisting on the PO approach, both, weak and strong nomological interpretations, are options worth considering for the interpretation of MBM’s WF. Neither fall prey to the problem of communication, nor put non-3D entities into the primitive ontology. Imposing these demands, a nomological reading of the WF seems to be the most promising way to classify MBM as a PO theory (while solving the problem of communication). Is this a conclusive argument to the effect that MBM’s WF must be read this way? Certainly not. For one, there might be other options beyond the ones presented in this paragraph. Second, the aim of the exercise was to show that MBM can be made compatible with the PO approach. It

can, but that, in itself, is merely responding to a demand, rather than justifying it. Those who deem it crucial that a theory be compatible with the PO approach may go for a nomological interpretation. Others may adopt a different one, dispositionalism_{MBM}, for example, which provides a solution to the problem of communication by endowing the particles with a real, non-3D property, but, because of that, isn't compatible with the PO approach. I will abstain from arguing in favour or against the alternatives: in my view, the interpretational openness of MBM is not a bug, but a feature.

Irrespective of this interpretational issue, MBM's metaphysical adequacy can still be questioned. The problem we are going to discuss in the following section harkens back to the roots of MBM, Bell's "Everett (?) theory", and has been dubbed "temporal solipsism". Challenging the latter's metaphysical adequacy, it is the reason for Bell's dismissal of the theory, and hence deserves some attention. At first glance, it stands to reason that MBM, like Bell's Everett (?) theory, is undermined by temporal solipsism. That this is not the case, we will see now.

II.5 A bit of history: Bell's Everett (?) theory and temporal solipsism

Despite its empirical equivalence with BM, and its solution to the problem of communication, MBM still faces the objection brought forward by Bell (1987b, 136): "Everett's replacement of the past by memories is a radical solipsism – extending to the temporal dimension the replacement of everything outside my head by my impressions." This shall be called "temporal solipsism". Doesn't temporal solipsism undermine the adequacy of MBM? After all, the present configuration is not connected to any past or future configuration. So it seems that any talk of my existence past or future is vacuous, a mere illusion, just like, according to the ordinary solipsist, the external world is nothing but illusory. To address this issue, let's have a quick look at the historical origin of MBM – Bell's "Everett (?) theory" – and why Bell rejected it immediately.

The first outline of a theory similar to MBM is to be found in the last section (5) of Bell's *Quantum mechanics for cosmologists*. There, Bell (ibid., 133ff.) presents what he calls "the Everett (?) theory", which "will simply be the pilot-wave theory without trajectories".²¹ He characterizes the Everett (?) theory as follows:

"Instantaneous classical configurations x are supposed to exist, and to be distributed in the comparison class of possible worlds with probability $|\psi|^2$. But no pairing of configurations at different times, as would be effected by the existence of trajectories, is supposed. And it is pointed out that no such continuity between present and past configurations is required by experience."

Here, Bell provides the general idea of a Bohmian theory without a GE. We are given the essential ingredients that make a theory Bohmian: possible configurations, i.e. configuration space ("the comparison class of possible worlds"), an instantaneous configuration x and the QEH for the distribution of ensembles of configurations as $|\psi|^2$. Despite not mentioning the GE – or rather, its omission – explicitly, he is eager to make clear that trajectories do not figure in the theory: configurations are taken to be instantaneous and "no pairing of configurations at different times, as would be effected by the existence of trajectories, is supposed."

After some remarks in which he compares his Everett (?) theory with the traditional account, Bell goes on to formulate a critique of the idea he just had introduced (136):

"In our interpretation of the Everett theory there is no association of the particular present with any particular past. And the essential claim is that this does not matter

²¹ To do actual history justice, it should be noted that, as Bell himself points out in a footnote, the origins of *Quantum mechanics for cosmologists* date back to an unpublished paper (*On the hypothesis that the Schroedinger equation is exact*, <http://cds.cern.ch/record/956196>) presented at Pennsylvania State University in September 1971, reproduced in *Epistemological Letters* in July 1978 and revised in *Speakable and Unspeakable in Quantum Mechanics* (1987a). Large parts of the final version were taken over from the original paper, including Bell's worries about temporal solipsism, although in the earlier version, the exact phrasing "radical solipsism" is not to be found.

at all. For we have no access to the past. We have only our ‘memories’ and ‘records’. But these memories and records are in fact *present* phenomena. [...] Everett’s replacement of the past by memories is a radical solipsism – extending to the temporal dimension the replacement of everything outside my head by my impressions. [...] If such a theory were taken seriously it would hardly be possible to take anything else seriously.”

After a quick comparison with creationism, Bell concludes his remarks on the Everett (?) theory. He quickly does away with, even ridicules the idea of a Bohmian theory without a GE (ibid.): “So much for the social implications. It is always interesting to find that solipsists and positivists, when they have children, have life insurance.” It can hardly be argued that his treatment is thorough, taking his MBM-like idea seriously.²² This is where Dürr and I disagree with Bell. We hold that MBM is not only empirically equivalent to QM – as discussed above and not contested by Bell – but that it is also metaphysically adequate. Temporal solipsism is not an issue. Showing the latter, i.e. showing that FAPP-continuous and consistent world histories can arise from a fundamentally stochastic, discontinuous micro dynamics, and explaining how this works, is the main objective of this thesis. In my presentation, I will expand on the presentation given in Dürr & Ehmman (ms). There, Dürr and I discuss Barbour’s proposal of time capsules, replacing the latter with a coarse-grained notion of worlds in terms of Everettian world branches and supplement it with an argument about their quasi-persistence to tackle the issue of temporal solipsism. While I agree with the in-principle viability of the time-capsule approach, I will critically evaluate Barbour’s construal of psychophysical parallelism, and the latter’s compatibility with the claim that individual, instantaneous configurations are sufficient for the experience of sentient beings. I present a detailed account of how FAPP-continuous, consistent world histories emerge from a

²² Barbour (1999, 300f.) agrees with this judgment about Bell’s premature dismissal of the theory: “This is all very entertaining – and I too have children and life insurance – but these are just the kind of *ad hominem* quips that were tossed at Copernicus and Galileo. I do believe that Bell came close to a viable cosmological interpretation of quantum mechanics, and should have kept faith with his title (‘Quantum mechanics for cosmologists’). But he left the cosmologists with nothing.”

random, discontinuous micro dynamics that surpasses the presentation in Dürr & Ehmman (ms). But before I can do that, we need to have a look at Barbour's approach.

II.6 Barbour: Temporal solipsism? Time capsules!

Amongst contemporary philosophers of physics, Julian Barbour is probably the one who takes Bell's idea more seriously than anyone else. However, the timeless cosmology he develops in his 1999 book *The End of Time* departs significantly from both Bell's Everett (?) theory and MBM. In his theory, there is no universal time. Everything there is, is eternal and static. In particular, configurations aren't actualized sequentially, one "at a time", as is the case in the Everett (?) theory and in MBM. In Barbour's cosmology, there is no time, so configurations, too, are eternal and static. There is no configurational change over time (ibid, 39): "If we could see the universe as it is, we should see that it is static. Nothing moves, nothing changes". Accordingly, experience of time has to arise from within such eternal, static configurations (ibid., 53): "The instant is not in time, time is in the instant." To account for the experience of time flow, of things moving, he therefore introduces the notion of time capsules (ibid.): "I shall outline, through the notion of time capsules, a theory of how a static universe can nevertheless appear to teem with motion and change."

Neither Dürr and I in our paper on MBM, nor I in this thesis, commit to Barbour's assumption of a fundamentally static universe, in the sense of a universe without fundamental, universal time. For MBM, we assume universal time. At each point in time, one configuration is actualized. As opposed to Barbour's cosmology, MBM's configurations are not eternal. Despite this important difference, Barbour's time-capsule approach is a candidate for tackling the issue of temporal solipsism. Barbour is well aware of Bell's Everett (?) theory and his criticism regarding temporal solipsism (cf. ibid., 300). In fact, he devotes an entire section of his book (299ff.) to the theory. Therein (299), he identifies Bell's notion of instantaneous

configurations containing records with his notion of time capsules: “[...] configuration points that can be called records. Although Bell did not use my term, such points are manifestly time capsules.” More generally, he introduces time capsules as follows (31): “Any static configuration that appears to contain mutually consistent records of processes that took place in a past in accordance with certain laws may be called a time capsule.”

One must think of such time capsules as configurations containing within them the complete history of a process up until that configuration, in the form of “snapshots”. Barbour (*ibid.*, 292) provides the following analogy: “Doting parents take daily snapshots of their child and stick them day by day into a progress book. The progress book after each successive day is like each successive point along the track in the big configuration space: it is the complete history of the child up to that date.” Other examples of records present in time capsules are geological formations (30) – “all geological formations, rock strata in particular, are now invariably interpreted by geologists as constituting a record (to be interpreted) of past geological processes” – or the particle tracks in a cloud chamber (299) – “alpha-particle tracks [...] have the obvious interpretation that they are records of alpha-particle motion” – all telling of a coherent history of events, apparently following some principles of time evolution that can be described in terms of laws of nature.

In addition to their structural richness and internal consistency, time capsules are supposed to be “sought out” by the WF (300): “[Bell’s] ‘Everettian’ interpretation is this: time exists, and the universal wave function Ψ evolves in it without ever collapsing. Because Ψ has the propensity to seek out²³ time capsules, it will generally be concentrated on them. Real events are actualized as follows. At each instant of time, Ψ associates a definite probability [...] with each configuration. At

²³ Barbour’s talk of the WF having “the propensity to seek out time capsules” is a bit obfuscating. (He repeatedly uses this metaphor, cf. Barbour 1999, 251, 296, 298, 300, 308, 316, 324.) What he means is that the probability density associated with the universal WF is concentrated over configurations that are time capsules, such that, at any time, the probability to be actualized is much higher for configurations that are time capsules than it is for configurations that are not time capsules.

any instant, just one event is actualized at random in accordance with its relative probability. The higher the probability, the greater the chance of actualization. Since time capsules have the highest probabilities, they will generally be selected.”

II.6.1 Explaining experiences

For his own, timeless cosmology, Barbour’s demand that time capsules are predominant is intertwined with “an assumption in the hallowed tradition of Boltzmann: *only the probable is experienced.*” (ibid. 266, his emphasis.) This assumption is controversial. When is a configuration probable enough to be experienced? Why, above all, should it be the case that only highly probable configurations are experienced at all? Barbour provides no good justification for his assumption. Nevertheless, at least at first sight, he has reason to make it despite: the vast majority of configurations in configuration space doesn’t admit of any experiences at all, and of those that do, the vast majority doesn’t admit of the reconstruction of FAPP-classical, coherent histories. Given that all configurations in his timeless cosmology are eternal, they are all actualized “at the same time” (which, in a timeless cosmology, can only be a figure of speech), which means that all possible experiences are actual. But if all possible experiences are actual, then by far most actual experiences are of incoherent histories. Yet, obviously, our own experiences are of the ordered kind – of consistent memories and records, admitting of the reconstruction of FAPP-classical, coherent histories. This seems unexpected. Our experiences, by all likelihood, should be incoherent.

Barbour’s assumption that the coherent time capsules are predominant, and that only the predominant configurations are experienced, allows him to reconcile the expectations derived from the theory with our experiences. *Prima facie*, this explains why we have the experiences we happen to have.

Given one demands this kind of explanation from Barbour’s timeless cosmology, one should also demand it from MBM. Other than Barbour, however, in order to

provide such an explanation, we don't have to assume that only the probable is experienced. On MBM, only one configuration is actual at a time. No matter how improbable its actualization is, once it is actualized, and contains experiencing beings, then these beings will have the experiences they just happen to have.²⁴ So when we demand for MBM that time capsules are assigned the highest probability to be actualized by the WF, what we want to achieve is that those configurations admitting of coherent histories are predominantly *actualized*, as opposed to being the only ones admitting of experiences in the first place. On MBM's sequential actualization dynamics, this is enough to reconcile theoretically derived expectations and our actual experiences. Barbour's controversial assumption isn't required.

Except for this difference, the rest of the reasoning goes quite similar for both theories. Time capsules should be assigned the highest probability to be actualized by the WF, because otherwise, e.g. assuming an equidistribution of probability over the totality of configuration space, almost all actualized configurations containing records at all would be configurations containing records that don't allow for the reconstruction of coherent histories, for such configurations vastly outnumber the rather special time capsules. As Barbour puts it (*ibid.*, 289): "among all possible worlds, the dull, disordered, incoherent states are overwhelmingly preponderant, while the ordered states form a miniscule fraction. But such states, sheer implausibility, must be presupposed if history is to be made manifest".²⁵ Of course,

²⁴ This is true under Barbour's further assumption that individual, internally static and eternal time capsules are sufficient for experiences to arise. A critique of this assumption will be spelled out later. For the present purposes, we can safely assume that it holds.

²⁵ The situation is reminiscent of the problem of Boltzmann brains. Given that the present state of our universe results from a random fluctuation out of equilibrium, we should expect ourselves to be Boltzmann brains: beings endowed with the necessary structure to hallucinate our experiences – including records, memories and anticipations of the future – but with not much more, for fluctuations into states that only contain such minimal beings hallucinating ordered structure and coherent records are much more likely than fluctuations into states where all the ordered structure and coherent records we seem to observe actually exist. (Cf. Carroll 2017.) Even worse: by far most Boltzmann brains fluctuating into existence should be "batty", such that the structures and records they hallucinate are disordered and incoherent. (Cf. Norton 2015.) Likewise, if MBM's actualization probability was equal for all configurations, we should expect most sentient beings to be Boltzmann brains, and most Boltzmann brains to be of the batty type. In fact, the analogy is quite to the point. Just replace the random fluctuations as supposed for the Boltzmann-brain scenario with random

by far the largest portion of configuration space consists of configurations that don't contain any records from which histories, coherent or not, could be reconstructed. The primary reason for demanding that time capsules are assigned the highest probability is not to ensure that in most actualizations, records are present, but to ensure that most actualizations in which records are present admit of a coherent history.

We must also revisit the issue of psychophysical parallelism: given strict dualism is not assumed, there has to be some “physical counterpart of psychological experience” (Barbour 1999, 302). Barbour follows Bell in his assumption that this is to be found in the configuration, i.e. in particle positions. From this assumption, he draws the following conclusion (ibid., 302, my emphasis): “it is clear from the way [Bell] makes memories and records responsible for our idea of the past, rejecting any ‘thread’ connecting configurations at different times, that subjective awareness of both positions and motions of objects must be derived from the structure in *one instantaneous* configuration. The self-sentient configurations must be time capsules. Not only the kingfisher but also the appearance of its flight must be *in one configuration*, for nothing else would be logically consistent”. Later on, I will contest this. It is not the case that subjective awareness of motions must arise from *one* configuration, and that nothing else would be logically consistent. I will argue that Barbour's claim contradicts our best understanding of the emergence of mental phenomena, and that he is aware of that: for all we know, the sentience of a sentient being arises not from one of its configurations, but from processes, involving change between several configurations. Barbour, pointing to the relevancy of brain processes (cf. ibid., 266), expresses a similar view. While he is right that upon “rejecting any ‘thread’ connecting configurations at different times, [...] subjective awareness of both positions and motions of objects must be derived

actualizations as featured in MBM. Given equal probability, of those randomly actualized configurations only few contain beings capable of experiences. And of those, in turn, almost all contain only batty Boltzmann brains. Were it not for the probability density to be distributed nonuniformly, we indeed should expect that amongst all brains instantiated, almost all are batty Boltzmann brains. Again, this is not logically ruled out by our experiences, but unexpected.

from the structure in *one instantaneous configuration*", I reject the antecedent for MBM: even assuming a stochastic, discontinuous micro dynamics as features in MBM, on a higher-level description, a "thread connecting configurations" can be had, namely in the form of world histories.

But first things first. For this introduction it shall suffice to assume that the criterion of psychophysical parallelism is met, i.e. that sentient beings arise from time capsules in such a way that a singular, instantaneous configuration is sufficient for their subjective experience. Given the probability density indeed concentrates on time capsules, we can expect that the experiences and memories of most actualized beings are of the regular, consistent kind, like ours are. Our history in its regularity, then, is nothing special.

But how is this predominant actualization of time capsules spelled out in MBM? Here comes into play the branching of the WF into FAPP-classical world branches, induced by decoherence, which accounts for the many-worlds character MBM inherits from the Everett interpretation.²⁶ With the WF being concentrated on FAPP-classical world branches, the actualization probability of configurations in the support of FAPP-classical world branches is overwhelmingly large compared to those in the support of non-classical branches. It is those configurations in the support of FAPP-classical world branches that contain mutually consistent records admitting of coherent histories, and that are called time capsules in Barbour's terminology. Now, since the actualization probability of configurations in MBM is given by the probability density $|\Psi|^2$, the super-particle, over the course of many actualizations, traces out FAPP-classical world branches, with overwhelmingly high probability. In other words, it almost always instantiates time capsules. We mustn't wonder why we find ourselves in such a "special", structurally rich configuration with consistent records, for, although such configurations might be special in the

²⁶ Everett's heritage is the many-worlds feature, not decoherence per se. The latter is by no means exclusive to Everett. Cf. Schlosshauer (2007, Ch. 8), discussing the role of decoherence for several (types of) "interpretations" (read: interpretations and alternative theories), or Bacciagaluppi (2016).

sense that they take up only a tiny volume in configuration space, their *actualization*, taking into account the branching structure of the WF and the resulting actualization probabilities, isn't special at all.

II.6.2 Fine-tuning for life and rare time capsules: an analogy

The arguments presented above made use of the probability density being concentrated on configurations admitting of coherent histories in order to explain the regularity of our experiences. However, it can be put into question whether such explanations are necessary and called for.

Let me try to flesh this out with an analogy between the rarity of time capsules and the fine-tuning for life of our universe. The parameters (e.g. constants of nature) defining our universe seem to be fine-tuned for life. Friederich (2017, 1.1.1) provides, amongst others, the following example of fine-tuned constants of nature: "The strength of gravity, when measured against the strength of electromagnetism, seems fine-tuned for life [...] If gravity had been absent or substantially weaker, galaxies, stars and planets would not have formed in the first place. Had it been only slightly weaker (and/or electromagnetism slightly stronger), main sequence stars such as the sun would have been significantly colder and would not explode in supernovae, which are the main source of many heavier elements [...]. If, in contrast, gravity had been slightly stronger, stars would have formed from smaller amounts of material, which would have meant that, inasmuch as still stable, they would have been much smaller and more short-lived".

Now consider the space of *possible* universes – a parameter space in which each relevant parameter is represented by a dimension. For example, the gravitational constant is represented by a dimension, and can vary over a range of different

values.²⁷ Likewise other universal constants. Each point in this parameter space represents a possible universe.

This parameter space, then, comprises a huge “number” of possible universes, all set up with different values for the relevant parameters. That only “few” (a small volume in parameter space, assuming there is an appropriate measure for that) of these possible universes admit of life like ours is the essence of the claim that our universe is fine-tuned (cf. Barnes 2012, 529). For our purposes of constructing an analogy, we may assume that this is indeed the case, as suggested by Friederich’s example. Under these conditions, if one actual universe would be selected at random from the vast space of possible universes, it would be highly improbable that this randomly selected universe is one of the few capable of harbouring life as we know it. This, as Barnes (ibid.) puts it, is what makes the claim that our universe is fine-tuned for life interesting: “it makes the existence of life in this universe appear to be something remarkable, something in need of explanation. The intuition here is that, if ours were the only universe, and if the causes that established the physics²⁸ of our universe were indifferent to whether it would evolve life, then the chances of hitting upon a life-permitting universe are very small.”

Standard explanations of fine-tuning for life are the deliberate creation of our universe by an intelligent designer, or the actual existence of a multitude of different universes. I am not going to engage in the debate about which to prefer. The point here is to make an analogy: on the one hand, a space of possible universes, only few of which admit of life; on the other, a space of possible configurations, only few of which admit of coherent histories. For MBM, we assume that this latter space is actually explored; the super-particle jumps through configuration space. Accordingly, to construct the analogy, we assume that the

²⁷ The terminology sounds a bit oxymoronic. How to imagine a *varying constant*? Of course, in this context, being a varying constant means being (effectively) constant in a given universe, but varying over parameter space.

²⁸ Barnes calls the parameters specifying a universe its “physics”.

former space is explored, too; a multitude of different universes from the space of possible universes is brought into being.²⁹

Given this multitude of universes is actual, it isn't surprising that the universe we inhabit, featuring just the right set of parameter values, exists. In the space of possible universes, there are universes featuring different sets of parameter values, most of them not allowing for any sentient beings at all. We certainly inhabit a special universe, with parameters suitable to us. It indeed is fine-tuned for our existence, but it is just one amongst the many possibilities explored. Further explanation for its fine-tuning seems called for only if we assume that there is only one universe. But if we assume many universes, as we do for this analogy, it is to be expected that some actual universe(s) will have the right conditions to bring about life like ours. Analogously, if we assume many configurations being actualized, as we do for MBM, no special concentration of the probability density over time capsules is necessary in order to explain our experiences. As it turns out, the arguments above, involving a special concentration of the probability density, attempt an explanation for something that is already explained.

Let's pause here for a moment. I have just pushed back against the idea that an explanation of our experiences must be given in terms of a special concentration of the probability density. But, in a sense, I still gave an explanation, albeit a different one, drawing on the fact that, on MBM, many configurations are actualized, one after another. One might wonder whether *any* explanation of us having the experiences we happen to have is needed *at all*.

II.6.3 Explaining the need for explaining experiences: a conjecture

Certainly, there isn't always a need to explain the occurrence of an unlikely event, like the actualization of the time capsule we inhabit, or the existence of the fine-

²⁹ As a side note, for the following argument, it doesn't matter whether these universes exist in parallel or one after the other.

tuned-for-life universe we inhabit (Friederich 2018, 2.2): “Even if fine-tuned conditions are improbable in some substantive sense, it might be wisest to regard them as primitive coincidences which we have to accept without resorting to such speculative responses as divine design or a multiverse. It is indeed uncontroversial that being improbable does not by itself automatically amount to requiring a theoretical response. For example, any specific sequence of outcomes in a long series of coin tosses has low initial probability (namely, 2^{-N} if the coin is fair, which approaches zero as the number N of tosses increases), but one would not reasonably regard any specific sequence of outcomes as calling for some theoretical response”.

Indeed, being surprised by the occurrence of an improbable event and then inferring that this event likely is one of a sufficiently long series of past events would be falling prey to the inverse gambler’s fallacy, originally described by Hacking (1987). Rowbottom (2015, IX.1) characterises it as follows: “[The inverse gambler’s fallacy] involves mistakenly inferring what happened in the *past* [...] because of what happens in the present. Or more precisely, in game-like scenarios, it involves the assumption that results at some point in time are *dependent* on earlier results in some way.” In other words, the (false) premise at work when committing the inverse gambler’s fallacy is

(IGF*) If an unlikely event occurs, then it likely is part of a long series of past events

Given that an unlikely event occurs, and IGF* is assumed, the (false) conclusion

(C*) The unlikely event likely is part of a long series of past events follows. In light of this, it seems best to abstain from any attempts at explaining unlikely events by inference to a past series of events.

That being said, the occurrence of unlikely events, at least psychologically, can be surprising, and people tend to feel the need for an explanation.³⁰ Let me try to

³⁰ There is some empirical evidence suggesting that in fact, people tend to commit the inverse gambler’s fallacy, see Oppenheimer & Monin 2009.

make sense of this apparent need for explanation, for the case of our universe's fine-tuning for life, and, by analogy, for our time capsules "fine-tuning" for the experience of a coherent history, exploiting Friederich's coin-tossing example.

In a series of ten tosses of an ideal, fair coin, I should not expect any *specific* sequence over any other *specific* sequence. The specific sequence HHHHHHHHHH, or the specific sequence HTHTHTHTHT, is just as likely as the specific sequence THTTHHHTTH (which I have just generated tossing an actual coin). For each individual coin toss, H and T both have probability $\frac{1}{2}$, independently of the other tosses, so each specific sequence of ten coin tosses has probability $\left(\frac{1}{2}\right)^{10} = \frac{1}{1024}$. They are all equally (un)likely. I'd commit the inverse gambler's fallacy if I inferred from a sequence of all heads that it likely is part of a long series of sequences.

However, flipping a coin ten times, as I did, I'd still be quite surprised by the outcome HHHHHHHHHH, while I wasn't surprised by THTTHHHTTH, the actual result I got. The former sequence triggers the (psychological) need for an explanation. Why is that so? Both specific sequences have the same probability. I really shouldn't be more surprised by one than by the other. One really shouldn't require an explanation while the other doesn't.

Here is my conjecture. The reason that I am surprised by HHHHHHHHHH but not by THTTHHHTTH might have to do with coarse-graining sample space into sets, resulting in a partition of sample space, in this case, into equivalence classes. Let me explain. Glossing over the specific order of Hs and Ts in sequences, i.e. only being concerned about the relative frequency of Hs and Ts, the latter sequence above is representative of a much larger set of sequences than the former. I'm surprised not because I got one specific sequence instead of another, equally likely sequence, but because I got a sequence from a tiny set instead of a sequence from a comparatively large set. (With all possible sequences being equally likely, the probability of ending up with a sequence from a certain set is the probability of

ending up with a specific sequence times the number of sequences in that set.) There are 252 sequences containing exactly five Hs. They form an equivalence class. No other equivalence class (with the equivalence relation “containing the same number of Hs”) in the sample space of possible sequences is larger. So having five Hs in ten tosses is the expected frequency. There are 210 sequences with exactly four Hs, and 210 with exactly six. 120 each with three or seven Hs, and so on. What's so surprising about the occurrence of the all-heads sequence, I figure, is that there is only one sequence in its equivalence class, while other equivalence classes are much bigger. The further the relative frequency of H's and T's deviates from the expected frequency, i.e. the smaller the relative size of the equivalence class the sequence is part of, the more surprised I am. My surprise is rooted not in the fact that a *specific* sequence occurred instead of another *specific* sequence, but in the fact that a sequence from the smallest equivalence class occurred instead of one from the bigger classes.³¹

I can imagine that this is why examples in presentations of the inverse gambler's fallacy usually are set up as they are. Hacking, e.g., construes an example with a gambler inferring from a roll of two dice showing double sixes that this roll most likely was part of a longer series of rolls. Why does Hacking construe the example with double sixes? Why not with a two and a four, say? It's just as likely a combination as two sixes are. However, an example with such a seemingly common combination would feel strange, wouldn't it? Why does it seem common, while double sixes don't? My hunch is that we're coarse-graining the sample space of possible outcomes into equivalence classes with respect to “being a double”. Then all the double somethings are put into a set, and all the rest into a different set, and since the former set is much smaller, the double somethings look surprising and in need of an explanation, motivating the unwarranted inference.

³¹ Even sequences with the expected frequency might psychologically be surprising. For example the sequence HTHTHTHTHT, or HHHHHHTTTTT. These look surprisingly ordered, as if they had been generated by applying a (deterministic) rule. So, to provide a tenable account for why some outcomes are psychologically surprising while others aren't is probably a bit more complicated than my simplified example suggests. However, for conveying the basic idea and making my point, it should be enough.

Let's look at another, even more illustrative example. Suppose a dart is randomly thrown once at a board with 32x32 equally sized squares. Only one of them is red, all other 1023 squares are green. Each square has the same probability of being hit as any other. But glossing over the fact that the squares are individual squares and partitioning the board under the equivalence relation "having the same colour" results in two equivalence classes, one much larger than the other. Suppose that the dart landed on the one red square. It isn't surprising that it landed on any *specific* of those 1024 squares. (If the squares all had different colours, or all the same colour, none would stand out.) But it is somewhat surprising – at least I would be surprised – that it landed on the single red square, and not on *any* of the 1023 green ones. This is what reeks of conspiracy.

Now, the question seems to be whether my surprise, on which my suspicion of conspiracy and demand for an explanation are based, is justified or not. Am I justifiably surprised that the dart landed on the red square? This, it seems, is a matter of perspective. Looking at the squares individually, on the "lower", not coarse-grained level of description, the situation presents itself as an equidistribution; all squares have the same probability of being hit. Because of that, being surprised by a specific square being hit isn't justified. On the higher, coarse-grained level of description, i.e. on the level of colours, probabilities aren't equidistributed; not all colours have the same probability of being hit. But if being surprised by an unlikely outcome, given there is a much more probable alternative (as opposed to many equally unlikely alternatives), is justified in general, then being surprised by the red square being hit, given all other squares are green, is also justified.

This puts me in an almost paradoxical situation. One and the same event, a specific square being hit, seems surprising or unsurprising, depending on whether I choose to coarse-grain for colour or not, i.e. on the partition I impose on sample space. If I do coarse-grain for colour, then I expect a green square to be hit almost always. The

red square being hit on a single throw thus triggers a need for explanation.³² If I do not coarse-grain for colour, I don't expect any outcome over any other outcome, and thus won't ask for an explanation. But with my expectations depending on coarse-graining, prospects of justifying my surprise seem forlorn. That is, unless the partition I impose itself is justified – if it is a “natural” partition, so to speak.

To get terminology out of the way, I call a partition “natural” iff it recovers actual, higher-level similarities and differences between the elements of sample space. To be more precise, a partition is natural iff it groups together those and only those elements that are equivalent with respect to an actual, higher-level property. For example, the partition of the dart board into colour equivalence classes is called natural because it recovers the actual, higher-level similarities and differences in colour of the squares: all and only red squares are grouped together in a set, all and only green squares are grouped together in another set.

You can see where this is going. Arguably, coarse-graining the parameter space of possible universes with respect to hospitability for life, or coarse-graining MBM's configuration space with respect to suitability to bring about the experience of a coherent history, result in natural partitions, too. Like the colour partition of the dart board, they recover actual, higher-level similarities and differences between the elements of these spaces. Like with the dart board, I shouldn't be surprised that any one of all the possible universes from parameter space was selected over any other. But being surprised that the one and only actual universe is hospitable to life seems to be justified.

³² A possible explanation for why the red square was hit, and not any of the green ones, is that the dart wasn't thrown randomly after all. It was Peter 'Snakebite' Wright, the current World Darts Champion, deliberately aiming for the red square. Another possible explanation is that the dart in fact was thrown randomly, but many times over. Throwing often enough, it will eventually land on the red square. In both these cases, I wouldn't be surprised by the red square being hit. Interestingly, both explanations involve me being mistaken about the initial setup: to make them work, I must revoke my assumption that the dart is thrown at random, or that it is thrown only once, respectively, and replace it with a different assumption. Note that assuming a proposition isn't the same as inferring it. In particular, assuming that the dart was thrown many times isn't to be mistaken for inferring that it was thrown many times.

So much for my conjecture about why certain outcomes seem surprising and psychologically demand for an explanation, while, when looking at the underlying sample space, all outcomes are equally likely. Even assuming the actualization probability density being flat over MBM's configuration space, the fact that my configuration was actualized does not demand for an explanation any more than the actualization of any other possible configuration. That I am still surprised by this actualization, and thus demand an explanation for it, is due to the partition I impose on configuration space: there are possible configurations, called time capsules, that allow for the experience of a coherent history, and there are other possible configurations that don't. Coarse-graining with respect to suitability to bring about the experience of a coherent history recovers these higher-level similarities and differences. The resulting partition is thus natural. It tells me that only a small set of possible configurations are time capsules, while all others aren't, and that time capsules thus are much less likely to be actualized than non-time capsules. Assuming there is only one configuration actualized, I'd feel very justified in being surprised that this one configuration was a time capsule, despite the fact that I know that each specific configuration is as likely to be actualized as any other (assuming actualization probability density being flat over configuration space).

Being justifiably surprised of, and thus demanding an explanation for, the actualization of a time capsule at some specific point in time, may I justifiably infer that, most likely, many configurations have been actualized before it, i.e. that the actualization of that time capsule most likely is part of a long series of actualizations, or would I fall prey to the inverse gambler's fallacy? Indeed, I would. Said time capsule could be the first configuration ever actualized, or the fifth, or the 10^{679} th. Luckily, however, I don't need to make this inference. The dynamics of MBM already tells me that uncountably many configurations are actualized in any finite time interval. This is built into the theory. I don't need to infer the likely existence of a series of previous actualizations, so the inverse gambler's fallacy is straightforwardly avoided.

II.6.4 Is MBM empirically incoherent?

On MBM's dynamics, the occasional actualization of time capsules is expected, even without a special probability density distribution. However, the decoherence-induced branching structure of the WF is relevant in a different matter. Barrett (1996) brings forward the charge of empirical incoherence against Bell's Everett (?) theory. He calls a theory empirically incoherent (ibid., 50) if it "fails to predict the existence of reliable records of an observer's measurement results to which the observer has epistemic access", and claims that the records present in time capsules are indeed unreliable, in the sense that the records contained within a given time capsule, e.g. our present one, typically are incompatible with records in previously actualized time capsules (ibid., 55): "the current particle configuration is independent of any past or future configurations [...] This means that one's records of measurement results would typically change in a pathological way over time and hence be wildly unreliable as records of what actually happened." He provides the following example:

"Suppose [...] that an observer M measures the x -spin of a spin-1/2 system S that is in a superposition of x -spins [...]. Suppose that the observer gets the result x -spin up for the outcome of her first measurement. Now what happens if the observer carefully repeats her measurement? The linear dynamics tells us that the quantum-mechanical state after the second measurement will be

$$\alpha|\uparrow, \uparrow\rangle_M|\uparrow\rangle_S + \beta|\downarrow, \downarrow\rangle_M|\downarrow\rangle_S$$

So what is the classical configuration after this measurement? [...] If M does in fact get 'x-spin down' for her second measurement, the classical configuration will now be one associated with the second term of the above state, which means that M 's 'record' of her first measurement will now read 'x-spin down', and it will thus appear, based on an examination of her records, that her two measurements did in fact yield the same result. More generally, one can show that the classical configuration would almost always be such that one's records would exhibit the

statistical correlations predicted by the standard theory whenever it makes unambiguous predictions.”

Barrett regards Bell’s Everett (?) theory – and a fortiori MBM – as empirically incoherent, because, at least prima facie, it’s records can be expected to be unreliable in this sense. However, he seems to overlook the import decoherence has on this matter. With decoherence, the WF branches into FAPP-classical world branches. Configurations are actualized at an arbitrarily high rate (see below), so consecutive actualizations of different time capsules within a world branch, because of the latter’s FAPP-classicality, can be expected to be mutually consistent with respect to the records they contain. So, contrary to Barrett’s fear, at least within a given branch, MBM’s records are epistemically accessible and reliable. They won’t be illusory in the sense of providing evidence contradicting what actually happened in the past. MBM thus satisfies Barrett’s requirements for empirical coherence. The charge of empirical incoherence is avoided.

11.7 Temporal solipsism solved?

This brings us back to the issue of temporal solipsism. Recall Bell’s complaint that “if such a theory were taken seriously it would hardly be possible to take anything else seriously”. One might still think that the theory is metaphysically deficient. After all, considering the vastness of configuration space, it is very unlikely that *precisely* the one time capsule I inhabit right now, this unique configuration, will be actualized again, even if the probability density is concentrated on time capsules, and consecutive time capsules within a world branch contain mutually compatible records. Consequently, I should exist for only one instant, such that all my memories of the past and anticipations of the future are rendered mere illusions. Even worse, not only my memories, but all records are meaningless – *bedeutungslos* in Frege’s sense³³: they point to a past that never existed. Annual

³³ Cf. Frege (1892).

rings in trees, the pictures of me as a child, geological formations including the fossils found in them, they all have been created at one instant, and cease to exist immediately after. In light of this, Bell likening the Everett (?) theory to Creationism seems understandable.

However, this problem arises only if time capsules are identified with individual, unique configurations, such that for any pair of actualized configurations $Q_{i,j}$, with $Q_i \neq Q_j$, different time capsules $TC_i \neq TC_j$ are instantiated, which is the case according to Barbour's presentation of Bell's Everett (?) theory (302): a time capsule is supposed to be "one instantaneous configuration". By contradistinction, "worlds", as mentioned earlier, are a coarse-grained concept. To provide a preliminary delineation: unlike time capsules, worlds do not consist of only one configuration. Instead, worlds are regions of finite volume in configuration space, comprised of all configurations instantiating macro states that are indistinguishable from a given macro state, i.e., worlds result from coarse-graining with respect to macroscopic features. The procedure will be explained in greater depth later.

Let's denote such a macro world \mathcal{W} , and the fact that one of its configurations is actualized at some time t by $(Q \in \mathcal{W}, t)$.³⁴ Since the set of configurations in \mathcal{W} is a finite region in configuration space, the probability of some configuration from \mathcal{W} being actualized at a specific $t = t_i$ depends on the size of that region and the probability density as given by the BR for that region and time:

$$P(Q \in \mathcal{W}, t_i) = \int_{\mathcal{W}} d^{3N} Q |\Psi_{t_i}(Q)|^2$$

Because worlds correspond to a region of *finite* volume in configuration space, this probability might be small, but is always non-zero (given the probability density over this volume isn't zero):

$$\forall t: 0 < P(Q \in \mathcal{W}, t)$$

³⁴ t , here, refers to universal time, i.e. the time parameter of the SE, describing the evolution of the universal WF.

Now take some finite time interval $I \subset \mathbb{R}$. With the probability of $Q \in \mathcal{W}$ being actualized at any arbitrary time t already being greater than zero, the probability of it being actualized within the time interval I , consisting of uncountably many instants, must be greater than zero as well: $P(Q \in \mathcal{W}, I) > 0$.

We can then select at random some T instants from I : $t_1, \dots, t_T \in I$. Think of T as the number of instants at which some configuration is actualized within the given time interval I , such that $R = \frac{T}{I}$ is the actualization rate (in actualizations per time span) for the time interval I . For example, if we select $T = 1000$ instants from a time interval I of one second, the actualization rate is $R = \frac{1000}{s}$. Note that $R \rightarrow \infty$ for $T \rightarrow \infty$ and finite time intervals. Because actualizations of configurations in general are stochastically independent – recall that $P(Q_t \cap Q_{t'}) = P(Q_t) * P(Q_{t'})$ – actualizations of configurations within a given world at different times are stochastically independent as well: $P(Q \in \mathcal{W}, t \cap Q \in \mathcal{W}, t') = P(Q \in \mathcal{W}, t) * P(Q \in \mathcal{W}, t')$. Despite this stochastic independence, we can expect that configurations from our world are actualized time and time again, contra Bell’s main worry: with the actualization probability $P(Q \in \mathcal{W}, t_i)$ for a configuration in \mathcal{W} to be actualized at t_i , the expectation value for some $Q \in \mathcal{W}$ to be actualized during I is the sum over the individual actualization probabilities at each instant t_i of the selected T instants $t_1, \dots, t_T \in I$:

$$E[Q \in \mathcal{W}; t_1, \dots, t_T] = \sum_i P(Q \in \mathcal{W}, t_i) = P(Q \in \mathcal{W}, t_1) + \dots + P(Q \in \mathcal{W}, t_T)$$

As can be seen, this expectation value depends on the number T of actualizations during I , i.e. on the actualization rate. If $R \rightarrow \infty$, $E \rightarrow \infty$ as well:

$$E[Q \in \mathcal{W}; t_1, \dots, t_T] \xrightarrow{R \rightarrow \infty} \infty$$

So, under the assumption that, during a finite time interval, arbitrarily many configurations are actualized, configurations from our world \mathcal{W} can be expected to be actualized arbitrarily many times in the same, finite time interval as well.

What’s the import of these considerations on our question at hand? Is Bell’s worry regarding the alleged temporal solipsism of MBM off the table? Indeed, as we have seen, on MBM’s construal of worlds, configurations within a world aren’t actualized

just once. In particular, configurations from the world inhabited by you and me are actualized infinitely many times. We can be reassured that they will be actualized again, and already have been actualized before, many times over. So, this particular worry is averted.

However, temporal solipsism can be understood as comprising more than the fear that any particular configuration is instantiated only once. Concerns may also arise about the temporal connection between different configurations at different instants of time. While Bell's original unease was particularly focussed on the "replacement of the past by memories" (Bell 1987b, 136), he at least hinted at a similar issue regarding the future, recall his comment about solipsists having life insurance (*ibid.*). His worries about the past being a mere illusion carry over to future configurations as well: in both directions of time, the connection between different configurations, in particular temporally *remote* ones, seems problematic. Recall that time capsules in Barbour's account are supposed to be static (Barbour 1999, 31): "Any static configuration that appears to contain mutually consistent records of processes that took place in a past in accordance with certain laws may be called a time capsule." But Barbour, by his claim that external, universal time does not exist, not only is forced to speak of *static* time capsules in the sense of *internally unchanging* configurations. Individual configurations are static in this sense by construal. Other than that, configurations are also supposed to be *static* in the sense of being "*not in time*" or "*eternal*". So there are no time capsules at *different times* in Barbour's account. Yet, assuming for MBM, by contrast to Barbour's timeless cosmology, that there are different configurations actualized at different times, there are temporally remote worlds, for example the world instantiated right now and the temporally remote world where I have children and life insurance, and the even more remote world where my future family benefits from said insurance. Likewise, the past world where I, as a young boy, am sitting in my bedroom playing with Lego and video games, knowing nothing about Barbour's concept of time capsules or even his book, which hasn't been written at the time. The world I inhabit now contains records of this earlier one – memories,

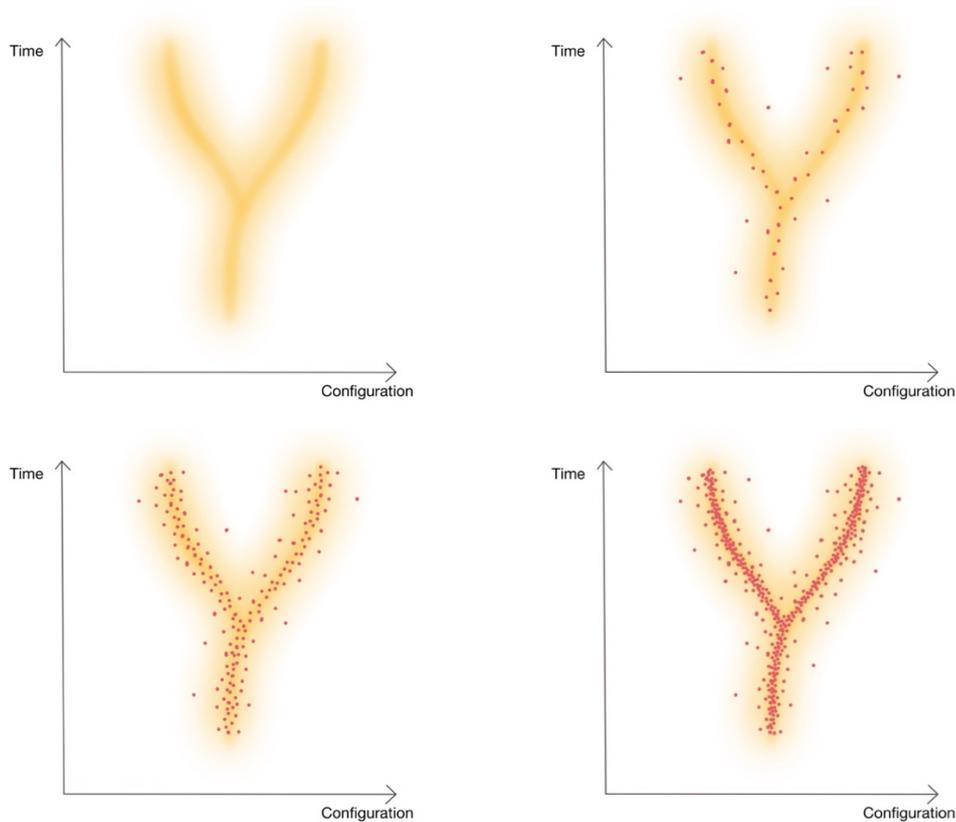
photographs, tape recordings. However, the relation is not symmetric. As Butterfield (2002, 319) puts it: “one time capsule can contain records of another, without the other similarly containing records of the first”. There was no record of my future self contained in the old world, just as there are no records of my children in the world I inhabit now. How are these instantiations of different worlds at different times linked up, over significant time spans in particular?³⁵ On their own, neither the fact that configurations from our world are actualized many times over, nor that consecutive time capsules within a world branch contain mutually consistent records, provides an answer to this question. In our paper, Dürr and I have tackled this issue by abandoning the original, Barbourian notion of static time capsules and resorting to Everettian world branches, which, for different instants of time, contain different, instantaneous configurations in their support, but themselves must not be considered instantaneous or static: they reach over extended time intervals, and hence comprise a history of configurations. And since world branches are a consequence of decoherence³⁶, this history is FAPP-classical, i.e. internally consistent. Because of that, records of the past, and likewise anticipations of the future, do point towards real instantiations of our world branch (Dürr & Ehmans ms, 35): “They permit inferences to our macro-world’s history; within each atto-second, infinitely many temporal fragments of it are realized. On

³⁵ Sticking with a timeless view, the question would have to be rephrased: How are worlds at different positions in configuration space linked up, such that they form a common history? For example, how is the world with my tea cup sitting on my desk related to the world with my tea cup on the floor, shattered into pieces? The answer Barbour would give is that the world with the broken cup contains records of the world with the unbroken cup – and likewise, for explaining how a sentient being can have the impression of perceiving the cup falling, records of the states in between. As we will see, the situation proves to be somewhat simpler for a model including universal time.

³⁶ Decoherence is a process that effectively suppresses interference effects via entanglement of a quantum system with its environment (Schlosshauer 2014, 1): “Realistic quantum systems are never completely isolated from their environment. When a quantum system interacts with its environment, it will in general become entangled with a large number of environmental degrees of freedom. This entanglement influences what we can locally observe upon measuring the system. In particular, quantum interference effects with respect to certain physical quantities (most notably, ‘classical’ quantities such as position) become effectively suppressed, making them prohibitively difficult to observe in most cases of practical interest. This is the process of decoherence, sometimes also called dynamical decoherence or environment-induced decoherence. Stated in general and interpretation-neutral terms, decoherence describes how entangling interactions with the environment influence the statistics of results of future measurements on the system.”

MBM, historical records aren't illusory or deceptive: They correctly describe a branch of reality (quite literally: an Everettian branch – that is, a quasi-classical history).” With world branches not only being instantiated many times over, but also forming quasi-classical histories, Bell’s temporal solipsism is warded off. While the fundamental picture still comprises independent, discontinuous actualizations of configurations, on a higher level of description, worlds branches (quasi-)persist.

Figure 1 provides an illustration of the quasi-persistence of world branches according to MBM:



*Figure 1.*³⁷ The universal WF provides the probability density (yellow) for actualizations of configurations (red). Actualizations happen in accordance with the probability density, the super-particle thereby carving out the WF in configuration space. The latter decoheres into world branches. For illustrative purposes, only two branches are depicted. Depending on the actualization rate R , configurations form a more or less dense history. From left to right, increasing actualization rates are shown. In the limit of $R \rightarrow \infty$, every world branch is instantiated infinitely many times in any finite

³⁷ This figure first appeared in Dürr & Ehmman (ms).

time interval. For suitable combinations of temporal coarse-graining and actualization rate, they quasi-persist.

This approach needs to be spelled out in greater detail than could be done in Dürr's and my paper on MBM. Combining elements of Barbour's time-capsule approach with elements from the Everettian theory, worlds correspond to finite volumes in configuration space. They are sets of (possible) configurations all of which (counterfactually) instantiate a macro state that is indistinguishable from a given macro state. Worlds, in some respect, resemble Barbourian time capsules, yet are importantly different: while each time capsule is a single configuration, finite volumes in configuration space are sets containing infinitely many configurations. In this respect, i.e. regarding their extension through configuration space, they more closely resemble Everettian world branches than Barbourian time capsules. However, they also differ from the former: Everettian branches do subsume histories of macroscopic change. Not so worlds. They, in themselves, are not histories of macroscopic change, since they only contain configurations that are macroscopically indistinguishable from a given macro state. In this respect, worlds are more like internally static time capsules. Nevertheless, we will see in a later chapter how world histories, i.e. successions of worlds, collectively instantiating macroscopic change, are formed from individual, internally unchanging worlds. Spelling out in greater detail how worlds are connected over extended time spans promises to be rewarding in several respects. For one, it will be beneficial to MBM and its acceptance as a viable alternative to BM and other non-relativistic quantum theories in that it provides a deeper, more thorough understanding of how world histories are formed from MBM's fundamentally discontinuous dynamics of randomly actualized configurations. For another, the kinship of worlds and time capsules, both being internally (quasi-)³⁸static, promises greater compatibility with Barbour's theory.³⁹ At the same time, the investigation will help to formulate a

³⁸ The qualification 'quasi-' applies to the notion of worlds, since they do comprise different configurations, such that actualizing one configuration within a world and then another within the same world amounts to microscopic change, while, at the same time, no macroscopic change occurs. By contrast, time capsules are individual configurations, which renders them static simpliciter.

³⁹ At least, I see no a priori reason for Barbour to not coarse-grain configuration space into worlds.

constructive criticism of the latter, possibly contributing to its improvement. I think there is good reason to be critical of Barbour's claim that the perception of macroscopic change can emerge from one individual, static configuration. In fact, based on our best scientific understanding, one may doubt that any kind of experience of sentient beings is possible from within only one such configuration. The presumption is that, in order to generate experience, in particular the perception of change, macroscopically relevant physical change is necessary, at least for beings the psychological states of which are instantiated by physical processes. Of course, there is no conclusive argument providing proof that we are such beings. But why throw overboard our best biology and neuroscience if we can have the cake and eat it, too? The suggestion I am going to put forward might solve this issue: we can have both, internally quasi-static worlds and the relevant macroscopic change as well, through the emergence of world histories.

The question that needs to be answered in order to make this proposal work is the following: how can world histories comprising macroscopic change arise from discontinuous instantiations of individual, internally quasi-static worlds? To tackle this problem, I will now delve into a related issue within Boltzmannian statistical mechanics (BSM). There, the standard way of partitioning phase space into disjoint macro regions corresponding to different macro states has unexpected consequences: upon traversing boundaries between macro regions, the continuous evolution of the micro state leads to a discontinuous change of the macro properties instantiated by the system in question. In fact, within the standard framework of BSM, whenever macroscopic change is modelled, it is modelled as discontinuous. This is at odds with the expectation that the thermodynamic evolution of a system, if not always, at least sometimes is continuous, and that BSM's models should be able to reflect that. This, I call the discontinuity problem. In the following chapter, I will make explicit the conditions that apply to the usual way of partitioning phase space. After exposing the discontinuity problem in detail and presenting three arguments against sharp boundaries between macro regions in phase space, I discuss some possible solutions to the issue at hand. Amongst them,

the conceptual shift from disjoint macro regions to overlapping worlds defined via a similarity relation shall turn out to be fruitful. The account thus developed – apart from some minor adaptations – almost straightforwardly carries over to the problem of the emergence of world histories in MBM, thereby hinting at a deeper connection between (the revised framework of) BSM and MBM, rooting from the shared concept of worlds and the coarse-graining relation between the realms of the micro- and macroscopic.

III. Smoothing out Boltzmannian statistical mechanics: On the viability of an alternative partitioning of phase space into overlapping worlds

A core concept in the usual framework of Boltzmannian statistical mechanics is the partitioning of phase space into non-overlapping macro regions, i.e. disjoint sets of micro states. Macro regions are of finite, non-zero volume, and stand in a one-to-one correspondence with different macro states.⁴⁰ Given a finite volume of phase space to be partitioned, there is a finite number of macro regions in a partition. Macro states are specified via particular macro values, e.g. for temperature, volume and pressure of a gas. In general, macro states are defined and specified via particular macro values taken on by macro variables: Following Werndl & Frigg (2015b, 1225), a system “can be characterized by a set $\{v_1, \dots, v_k\}$ of *macrovariables* ($k \in \mathbb{N}$). The v_i assume values in [macro value space] \mathbb{V}_i , and capital letters V_i denote the values of v_i . A particular set of values $\{V_1, \dots, V_k\}$ defines a *macrostate* M_{V_1, \dots, V_k} .”

Boundaries between macro regions are construed as being sharp: under a given partition, each micro state is a member of one and only one macro region. Since macro regions stand in a one-to-one correspondence with macro states – that is,

⁴⁰ Cf., e.g., Callender (1999), Frigg & Werndl (2012), Werndl & Frigg (2015b), Goldstein et al. (2017).

there is a bijection between the set of macro regions and the set of macro states – each micro state instantiates one and only one macro state, and micro states from different macro regions instantiate different macro states. Also, once fixed, a given partition doesn't change (Hemmo & Shenker 2012, 53 and 56): “The partition of the state space into macrostate regions is fixed in time”; “the partition of the state space into macrostates [...] is time-independent”. This partition of phase space into a finite number of macro regions with sharp, fixed boundaries is conceptually puzzling, for it results in models that are unable to describe continuous macroscopic change: whenever macroscopic change is modelled within the framework of BSM, it is discontinuous. BSM's models predict that every time the micro state of a system crosses the boundary between two macro regions, a discontinuity in the system's macro values results. However, models that describe macroscopic change as always discontinuous are at odds with our expectations: a system's macroscopic evolution as described by thermodynamics involves continuous change of its macro properties – of the temperature value, for example – so we wouldn't expect to encounter discontinuities whenever a system's macro properties change.

Likewise, continuous evolution of the micro state should not lead to discontinuous change of the macro state, i.e., should not lead to a discontinuity in the evolution of macro values. Yet, as the framework of BSM is set up, continuous change of the micro state sometimes (i.e. upon traversing boundaries between macro regions) *does* lead to discontinuous change of the macro state. In fact, the models resulting from the framework of BSM predict that, *whenever* there is macroscopic change, it is sudden and of a jumpy nature: during its wanderings through a macro region, the micro state always and invariably instantiates the same, unchanging macro state, according to the framework's models. As soon as the micro state moves into a different macro region, however, it suddenly instantiates a different macro state: since boundaries between macro regions are construed as sharp, macro values jump upon traversal of the micro state from one macro region to another. By contrast, since macro regions and macro states stand in a one-to-one correspondence, and macro states, in turn, are specified by a set of macro variables

taking on constant macro values, the latter remain completely unchanged as long as the micro state travels *within* a macro region. This results from the definition of macro states via “particular set[s] of [macro] values”, and the bijective association of the thusly defined macro states with macro regions of finite volume (Werndl & Frigg 2015b, 1225): “a system’s microstate uniquely determines its macrostate. Every macrostate M is associated with a macroregion Γ_M consisting of all $x \in \Gamma_E$ [with Γ_E denoting the energy hypersurface, i.e. the complete available phase space for a constant energy] for which the system is in M .” The consequence of this construal is that macro values always jump upon the micro state crossing boundaries between macro regions. Assuming that the thermodynamic evolution of macro states is to be recovered, jumping macro values, especially upon tiny changes of the micro state, seem inadequate. Yet, the partition of phase space into disjoint macro regions of finite, non-zero volume with fixed, sharp boundaries, as prescribed by the usual framework, results in models describing jumping macro states.

This issue, merely sketched so far, I call the discontinuity problem (DP) of the standard framework of BSM. It will be discussed in detail in the following. At this point, I should probably put some emphasis on the fact that DP only pertains to models resulting from the usual framework of BSM, and explicitly not to the target systems described by those models. To be perfectly clear: I am not claiming that the physical systems themselves suffer from DP, i.e., that they always undergo discontinuous macroscopic change, if they undergo macroscopic change at all. Rather, only the models that describe them entail these unavoidable discontinuities. This, in itself, would even be unproblematic. Per se, there is nothing wrong about models allowing only for discontinuous change. DP arises only if we want to have these models adequately describe systems that undergo continuous macroscopic change. So, what DP actually refers to, is the discrepancy between what we want or expect from these models, and what they are able to deliver.⁴¹

⁴¹ One might be tempted to put into question whether DP really poses a problem: aren’t most (if not all) models deficient in some or other respect? If so, why think that the deficiency of BSM’s models is

In what follows, I will first give a quick overview of ‘The usual presentation of BSM’. Next, I will make some ‘Remarks on macro regions and partitioning phase space’, wherein I describe and formalize the central conditions for the usual way of partitioning phase space, i.e. the *no-overlap* condition and the *exhaustive-cover* condition. I then describe in some detail ‘The discontinuity problem as it arises from the framework of BSM’ and discuss ‘BSM as a fundamental theory in light of the discontinuity problem’. Afterwards, I present some ‘Hints of digression from the usual’ as found in the recent literature. In ‘Against sharp boundaries of macro regions in phase space’ (I), (II) and (III), I provide three arguments making explicit the consequences resulting from DP. Equipped with a proper understanding of the problem, I will discuss two possible solutions in ‘Towards an alternative’: ‘Removing exhaustive-cover’ and ‘Removing no-overlap’. Both shall turn out to be unsatisfactory. For ‘A different solution’, I then suggest making a conceptual switch ‘From macro regions to worlds’. The latter shall be defined by a relation of *sufficient similarity*. In ‘Introducing a new framework for BSM: Some conditions’, the concept shall be explicated and formalized. I will discuss ‘Some anticipated worries’ regarding the new framework and show how it solves DP as encountered within the old framework by going ‘Back to the example’ used before. Finally, I will discuss issues regarding ‘World size, the equilibrium world, and entropy increase’ before closing with the ‘Conclusion’ of the chapter.

III.1 The usual presentation of BSM

Let’s start our investigation of the discontinuity problem sketched above by taking a closer look at the framework of BSM as it is usually presented, e.g. in discussions of approach to equilibrium:⁴² We are interested in a (classical) system comprised of N constituents that move through physical space. The micro state of the system – a

problematic? I will respond to this after introducing DP. I am going to argue that DP is particularly pressing if one regards BSM as a fundamental theory, as Frigg & Werndl (2019) do.

⁴² Roughly following the presentation in Frigg & Werndl (2012, 101f.).

complete specification of the positions and momenta of all the constituents involved in the system – is represented by a point x in its $6N$ -dimensional phase space Γ . If the energy of the system is constant, the phase space volume available to the system is reduced to a $(6N - 1)$ -dimensional energy hypersurface Γ_E . Sub-regions within Γ_E are specified such that they correspond to different macro states M_i – specified by particular macro values taken on by macro variables, see above – of the system: The sub-region in Γ_E corresponding to a macro state M_i is the set of all $x \in \Gamma_E$ that instantiate M_i and called the macro region Γ_{M_i} . Macro regions Γ_{M_i} form a partition of Γ_E : they cover all of Γ_E and do not overlap. A macro state of special interest is M_{eq} , the equilibrium state, instantiated by any micro state $x \in \Gamma_{M_{eq}}$ (and counterfactually instantiated⁴³ by all micro states $x \in \Gamma_{M_{eq}}$).⁴⁴ Given the standard Lebesgue measure, the volume of sub-regions of Γ_E can be assessed and hence is rendered comparable. It is usually said that (Frigg & Werndl 2012, 102) “for gases $\Gamma_{M_{eq}}$ is vastly larger [...] than any other macro-region, a fact also known as the ‘dominance of the equilibrium macrostate’ [...]; in fact Γ_E is almost entirely taken up by equilibrium microstates. For this reason the equilibrium state has maximum entropy.”

III.2 Remarks on macro regions and partitioning phase space

⁴³ “Counterfactually instantiated” in the sense that at any time, only one micro state is actualized, and thus only one micro state actually instantiates the macro state, but the actualization of the other micro states in the same macro region *would* instantiate the same macro state, if they *were* actualized instead.

⁴⁴ Answering the question (Werndl & Frigg 2015a, 19f.) “under what circumstances [...] an equilibrium state exists?” is not as simple as the usual mantra, according to which the equilibrium state is simply the macro state corresponding to the largest macro region, might suggest (ibid., 26, their emphasis): “for an equilibrium to exist *three factors need to cooperate: the choice of macro-variables, the dynamics of the system, and the choice of the effective state space* Z [e.g. the $(6N-1)$ -dimensional energy hypersurface, as opposed to the *full state space*, i.e. $6N$ -dimensional phase space (my remark)]. The cooperation between these factors can take different forms and there is more than one constellation that can lead to the existence of an equilibrium state. The important point is that the answer to the question of existence is holistic: it not only depends on three factors rather than one, but also on the interplay between these factors.” Werndl & Frigg (2015a) provide detailed arguments, to which the reader is referred. For our purposes, it is safe to assume that an equilibrium macro state exists, for even if not, the issue of macroscopic change upon the micro state crossing boundaries between macro regions remains.

Of special importance for our present purposes is the notion of macro regions and their partitioning phase space (ibid.): “From a macroscopic perspective, the system is characterised by a set of *macrostates* $M_i, i = 1, \dots, m$. To each macrostate corresponds a macro-region Γ_{M_i} consisting of all $x \in \Gamma_E$ for which the system is in M_i . The Γ_{M_i} form a partition of Γ_E , meaning that they do not overlap and jointly cover Γ_E .” Frigg et al. (2016, 4.1) summarize the key feature of the correspondence between micro states and macro states in a short statement: “to every given microstate x there corresponds *exactly one* macrostate”.

For one thing, this means that x either instantiates a certain macro state M_i or another macro state M_j , but never both. Since macro states M_i and M_j correspond to macro regions Γ_{M_i} and Γ_{M_j} , respectively, x must not belong to two macro regions. This we call the *no-overlap* condition:

$$\forall x \forall \Gamma_{M_i, j} \neg (x \in \Gamma_{M_i} \wedge x \in \Gamma_{M_j}); i \neq j \quad (\text{no-overlap})$$

For another, it means that macro regions exhaustively cover the available phase-space volume Γ_E . All micro states must instantiate *exactly one* macro state, read as: every micro state must instantiate *some* macro state, and hence belong to *some* macro region. It cannot belong to *no* macro region at all. Hence, all of Γ_E must be covered by macro regions (“they ... jointly cover Γ_E .” (Frigg & Werndl 2012, 102.)) Let’s call this the *exhaustive-cover* condition:

$$\forall x \exists \Gamma_{M_i} (x \in \Gamma_{M_i}) \quad (\text{exhaustive-cover})$$

Taken together, *no-overlap* and *exhaustive-cover* ensure that there is always some macro state instantiated, that this macro state is internally consistent, i.e. that a micro state never instantiates two (or more) distinct macro states at the same time, and that macro regions are not fuzzy: while it might be the case that there is some uncertainty or even complete ignorance about which macro state a given system is

actually in – if only because the respective macro values haven't been measured/observed – we can always be sure that its actual micro state x is located in *some* macro region or other, i.e. in one (*exhaustive-cover*) and only one (*no-overlap*) macro region Γ_{M_i} , and that thus, one and only one macro state is instantiated at any given time. After all, fuzzy boundaries would mean that the actual micro state does not always belong to a certain (i.e. to one and only one) macro region. Being located at the fuzzy boundary, x could belong to Γ_{M_i} , to Γ_{M_j} , to neither, or to both.⁴⁵ The latter option “both” is excluded by *no-overlap*. The option “neither” is excluded by *exhaustive-cover*. Hence, it must belong to either one of the former. In other words: boundaries of macro regions are sharp.

These considerations about the overlap and fuzziness of macro regions – or rather: these stipulations of their impossibility – are reflected in the schematic pictorial representations of phase-space partitioning we find, e.g., in Goldstein et al. (2019)⁴⁶:

⁴⁵ For an illustration of fuzziness, think of the paradigmatic example of the fuzziness of clouds. Clouds can be said to be fuzzy: it is not always clear whether a certain water molecule belongs to a certain cloud, to another cloud, to neither, or to both. The – quite literally – fuzzy nature of clouds renders it difficult to provide a clear and uncontroversial answer.

⁴⁶ I thank Sheldon Goldstein and the publisher of *Annalen der Physik* for the permission to reproduce this figure. It first appeared in Goldstein et al. (2017). Copyright Wiley-VCH Verlag GmbH & Co. KGaA. Reproduced with permission.

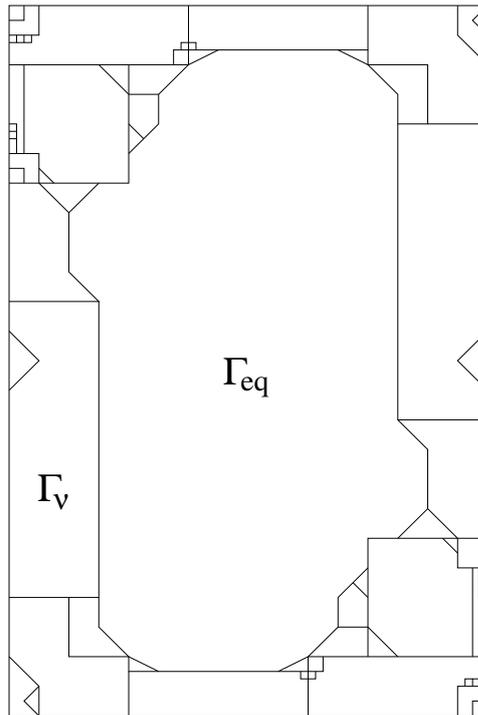


Figure 2

As can be seen, macro regions cover all of phase space, do not overlap and have sharp boundaries.

III.3 The discontinuity problem as it arises from the framework of BSM

It has been stated before (see above) that macro regions in phase space are of finite, non-zero volume, and that there is a finite number of macro regions in a partition. Why is that so? First, let's answer the question why macro regions must have a finite, non-zero volume. The reason for this lies in the way BSM works. The Boltzmann entropy function $-S_B(M_i) := k_B \log [\mu(\Gamma_{M_i})]$ ⁴⁷ – relies on the volume of macro regions in order to assign different entropy values to different macro states of a system. Recall, for example, what (Frigg & Werndl 2012, 102) say about the entropy value of the equilibrium state and the volume of the corresponding macro region: “for gases $\Gamma_{M_{eq}}$ is vastly larger [...] than any other macro-region, a fact also known as the ‘dominance of the equilibrium macrostate’ [...]; in fact Γ_E is almost

⁴⁷ Cf. Frigg & Werndl (2012, 102), Werndl & Frigg (2015b, 1225).

entirely taken up by equilibrium microstates. For this reason the equilibrium state has maximum entropy.” Now assume, contrary to this, that macro regions in fact were of volume zero. The entropy function then would assign to macro states no well-defined entropy value. A theory allowing for such models would hardly be able to accomplish the things BSM does, and certainly wouldn’t deserve the name BSM. For example, arguments about the relation of entropy change and the arrow of time would be completely blocked. This consequence, especially for the purpose of discussing the framework of BSM as it stands, is unacceptable. So indeed, it seems safe to say, in accordance with the way BSM is standardly presented, that macro regions in BSM are construed as having finite volume greater zero.

Let’s move on to the second part of the statement above: there is a finite number of macro regions in a partition. The reason for that is rather simple. As we have just seen, macro regions, as construed within the framework of BSM, are of finite, non-zero volume. It is trivial that the finite volume of phase space available to a system can only be partitioned into a finite number of compartments – macro regions – given the latter are of finite, non-zero volume (and do not approach zero volume). But why, one might want to ask, do we even assume that the available phase space volume is finite? If it would be infinite itself, we’d be able to partition it into an infinite number of compartments, despite the latter having a finite, non-zero volume. However, it is unreasonable to assume an infinite phase space volume for at least many of the systems BSM wants to model. Think, for example, of the paradigmatic gas in a box. Clearly, the spatial degrees of freedom each constituent has are bounded: particles can only move within the box (itself of finite volume). BSM wants to be able to model such cases, so it is uncalled for to simply assume that BSM only models systems that aren’t limited to a finite spatial volume. Likewise, for finite temperatures, the momenta of the particles can only take on values within certain limits. Accordingly, the phase space volume of the entire system is restricted to be finite, if very large. And once we have a finite volume of phase space, it is impossible to partition it into an infinite number of compartments with finite, non-zero volume. So the number of macro regions must be finite.

Now that we have established that there is a finite number of macro regions, we can already see that, in models of BSM, there can only be a finite number of different macro states. Macro states and macro regions, as I have said before, stand in a one-to-one correspondence. That is, the mapping between macro regions and macro states is bijective. But it is impossible to map a set containing an infinity of elements – let alone a set containing an uncountable infinity of elements – bijectively onto a set containing a finite number of elements. Accordingly, it is impossible to bijectively map a set containing infinitely many macro states onto a set containing finitely many macro regions. So if we have a finite number of macro regions, we are forced to let them correspond to only a finite number of different macro states. In fact, the number of macro states and the number of macro regions in the model must be equal: there are precisely as many macro states as there are macro regions, and this number is finite.

Having established that there is a finite number of different macro states, we will now find that there is a finite number of different sets of macro values in the models resulting from the usual framework of BSM. Macro states, as defined above, are specified by sets of macro values. The simplest case, of course, is if this set of macro values specifying a macro state only contains one relevant element, such that we can ignore the rest. Take, for example, some macro states that differ only in their temperature macro values. In this case, we can ignore other macro variables: they are assumed to take on identical values for the different macro states. Since there is a finite number of different macro states in the model, and these macro states differ in their temperature macro value, there is a finite number of different temperature macro values. The reason for this is similar to the one we just discussed in the case of macro regions and macro states: the set of different macro states, since the latter are specified by temperature macro values, stands in a one-to-one correspondence to the set of different temperature macro values. So with a finite number of different macro states, there is a finite number of different temperature macro values specifying the former.

Now, it is important to recognize that *different* macro regions must correspond to *different* macro states, and thus to *different* macro values. For if different macro regions would correspond to the same macro state, i.e. to identical macro values, they wouldn't be different macro regions in the first place. Such macro regions, corresponding to the same macro state, would be combined into one macro region. This is a consequence of the one-to-one correspondence between macro regions and macro states in the framework of BSM. In effect, different macro regions must correspond to different macro states, and, a fortiori, to different (sets of) macro values. In the simplified example, where we are only concerned about the temperature macro value, i.e. assume that all other macro values are identical for the different macro states, this means that different macro regions must correspond to different temperature macro values. But as we have established, there is only a finite set of different temperature macro values. What, then, does it mean for a finite set of temperature macro values to consist of different temperature macro values? It means that, if we take the difference between any two of the values taken from this set, this difference must be non-zero. That is to say: take any pair of macro values from the finite set of macro values, and subtract one from the other. The result must always be different from zero, for if it was zero, the two values would be identical. (They also can't arbitrarily approach each other, because the set of macro values is not a continuum.) It is unreasonable to assign identical macro states, i.e. macro states not differing in their macro values, to different macro regions, as argued before: those macro regions would be combined into one macro region, such that, in the end, macro regions and macro states remain in a one-to-one correspondence, and likewise macro regions and (sets) of macro values.

So what we are left with is a finite set of different macro values, each specifying a different macro state, corresponding to different macro regions. But if all pairs of elements of the set of macro values are different in that the subtraction of any one of them from any other results in a non-zero difference, then the transition

between two macro values, i.e. between two macro states, i.e. between two macro regions, is always modelled as discontinuous: any non-zero interval on the reals contains an uncountable infinity of real numbers. And since any two macro values from the set of macro values assigned to macro states, and thus to macro regions, must differ, they must be separated by a non-zero interval, which can't arbitrarily approach zero, because the number of different macro values must be finite. So the transition between any two macro values from this finite set of macro values must be discontinuous. A fortiori, the transition between any two macro states must be discontinuous, and accordingly, the transition between any two macro regions results in a discontinuity. So we can say that the transition of the micro state from one macro region to another, corresponding to the instantiation of different macro states, and thus of different (sets of) macro values, always results in a discontinuity (in the standard framework of BSM). To illustrate this, imagine the temperature macro value of a system as modelled by BSM graphed as a function over time (figure 3):

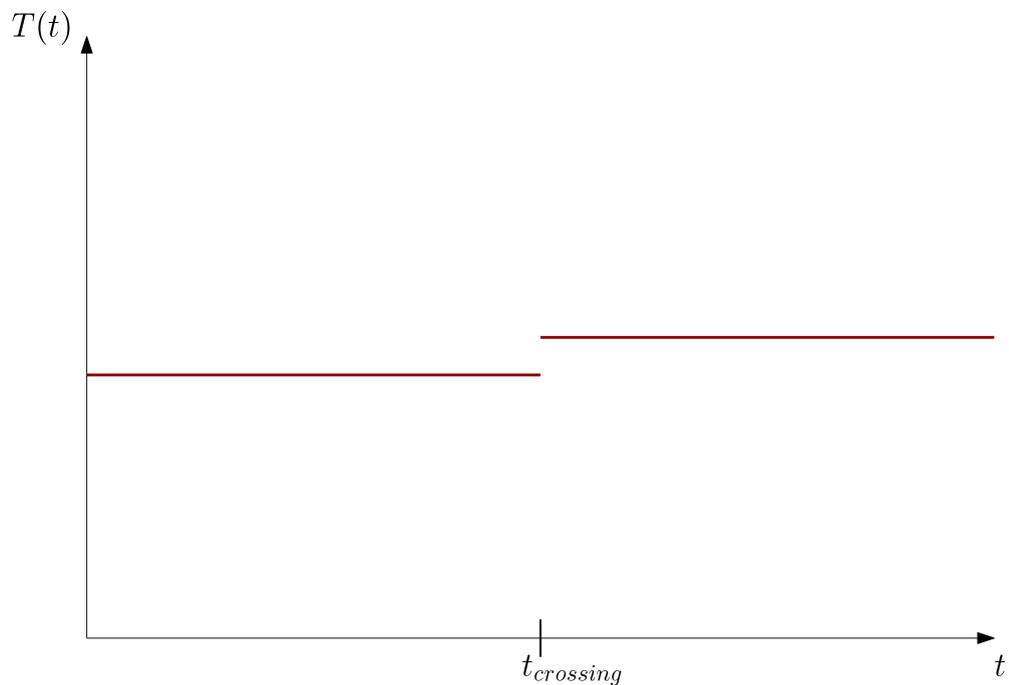


Figure 3

As long as the micro state wanders through a macro region Γ_{M_i} , the temperature macro value $T(t) = T_{M_i}$ remains absolutely constant, according to the model. So the graph shows a continuous, straight line, representing the constant temperature macro value. But at the precise instant the micro state crosses the boundary and moves into another macro region Γ_{M_j} – call this instant $t_{crossing}$ – the temperature macro value changes to $T(t) = T_{M_j}$. Then, as long as the micro state stays in this macro region, the temperature macro value again remains absolutely constant, so the graph again shows a continuous, straight line, representing this different temperature macro value. Thus, the graphed function, at $t_{crossing}$, contains a discontinuity. And this discontinuity is not removable, as is easily shown: consider the limit $L^- = \lim_{t \rightarrow t_{crossing}^-} T(t)$ of the function for t approaching $t_{crossing}$ from the left side, i.e. for the time during which the micro state moves through the first macro region Γ_{M_i} . The value of this limit is the temperature value before crossing the boundary: $L^- = T_{M_i}$. Now consider the limit $L^+ = \lim_{t \rightarrow t_{crossing}^+} T(t)$ of the function for t approaching $t_{crossing}$ from the right side, i.e. for the time during which the micro state moves through the second macro region Γ_{M_j} , traced backwards. The value of this limit is the temperature value after crossing the boundary, $L^+ = T_{M_j}$. Since $T_{M_i} \neq T_{M_j}$, the limit from the left side and the limit from the right side are unequal as well: $L^- \neq L^+$. So there exists no single limit. So the discontinuity at $t_{crossing}$ indeed is a jump discontinuity.

This is how the discontinuity problem arises from the usual framework of BSM. As I have said before, DP consists in the discrepancy between the expectation that thermodynamic macro values change continuously, and the fact that, in the usual framework of BSM, continuous change of macro values can't be modelled. So it neither consists solely in the fact that BSM is unable to model continuous change alone, nor does it contain the claim that in nature, or in thermodynamics, or in reality, macroscopic change is always discontinuous. The discontinuity we encounter in BSM's models is an artefact resulting from the way the framework is

set up: partitioning phase space into a finite number of macro regions, standing in a one-to-one correspondence with different macro states, the latter themselves standing in a one-to-one correspondence with different sets of macro values, such that, in the end, the elements of a finite set of macro regions stand in a one-to-one correspondence with the elements of a finite set of sets of macro values – – this partitioning, judged from the thermodynamic viewpoint, where continuous change of macro values is allowable, is an idealisation (see below). One might reasonably prefer a framework that is able to avoid it.

The discontinuity problem of BSM is put into sharp relief when regarding BSM as a fundamental theory, as Frigg and Werndl (2019) do. They differentiate between fundamental and effective theories of statistical mechanics (SM), classifying BSM as the former, and Gibbsian statistical mechanics (GSM) as the latter (*ibid.*, 425 and 435): “BSM is a fundamental theory while GSM is an effective theory. This means that BSM provides a true description of the systems within the scope of SM; GSM offers an algorithm to calculate values defined by the fundamental theory. The algorithm is often easier to handle than the fundamental theory and provides results where the fundamental theory is intractable. As every effective theory, GSM works only within a certain domain of application.” “Asking whether GSM provides a correct fundamental description of the world, or, if the answer to this question is negative, trying to revise GSM so that it does provide such a description, is a mistaken endeavour. Effective theories do not offer fundamental descriptions; they are calculatory devices of instrumental value; no more and no less.”

By contrast to effective theories, fundamental theories are supposed to provide correct descriptions of the world (*ibid.* 431): “BSM is quite unlike GSM [...]. Dynamical considerations occupy centre stage in BSM. It introduces macrostates with corresponding macroregions, and then defines equilibrium in explicitly dynamical terms (namely as the macrostate whose macroregion is such that, in the long run, the system’s state spends most of its time in that macroregion). [...] [U]nder the assumption that the world is governed by Newton’s equation of motion

[...] the dynamics considered in BSM is the true dynamics at the fundamental level: the unabridged and unidealised dynamics with all interactions between all microconstituents of the system. Equilibrium results from macrostates that are defined in terms of macrovariables that supervene on the true microdynamics of the system, and where a system fluctuates away from equilibrium it does so as a result of the true underlying dynamics. In a classical world the theory gives a full account of all this – nothing is left out and nothing is averaged over. BSM provides the complete fundamental theory of SM systems.”

III.4 BSM as a fundamental theory in light of the discontinuity problem

Above, I have promised to respond to the claim that DP isn't problematic after all. To this, I shall come now.

First, I need to distinguish between two types of functions models can have, and the way in which they can be deficient.

The first type of function of models is to “*represent* a selected part or aspect of the world, which is the model's target system.” (Frigg & Hartmann 2020, 1., my emphasis.) Scale models have this function. They are (ibid.) “down-sized or enlarged copies of their target systems. A typical example is a small wooden car that is put into a wind tunnel to explore the actual car's aerodynamic properties.” Whether a model fulfils its representational function successfully depends on the application. The wooden car model might be successful in representing the aerodynamic aspects of the real car, but it certainly fails in representing the real car's material composition (cf. ibid.). It is deficient in the sense that it doesn't represent all aspects of the real car. With respect to the intended application, this deficiency isn't problematic. There, its function is to represent a certain aspect of the real car, not all of them.

A different type of function of models is to *describe* a target system that falls within the domain of description of a given theory or theoretical framework. For example, a system of two massive objects orbiting each other in empty space is a target system that falls within the domain of description of Newtonian classical mechanics. Models describing a target system often entail approximations or idealisations. Borrowing the distinction from Norton (2012, 210), “an approximation is an inexact description of a target system.” For example, using Newtonian mechanics to describe the target system of two objects orbiting each other in empty space results in a model ignoring relativistic effects. This amounts to an approximation. “An idealization”, on the other hand, “is a real or fictitious system, distinct from the target system, some of whose properties provide an inexact description of some aspects of the target system.” (Fictitiously) modelling the above two-body target system as a system of two point-like masses would be an idealization. This model isn’t merely an approximation. It also has a representational function in that it represents the aspects “mass” and “position” of the target system in a fictitious system of point-like masses. Different functions of models aren’t necessarily mutually exclusive. Approximations are deficient in virtue of being inexact, and the same is true for idealisations, for they entail approximations. Idealisations however, because they are systems themselves (if only fictitious ones), may also be deficient in virtue of not representing all aspects of the target system. Again, these deficiencies aren’t problematic, as long as the model fulfils its intended function.

The function of BSM’s models is of the second type. They describe some aspects of a target system, e.g. of a gas in a box, that falls within the domain of description of BSM. To be more precise, they describe the macro evolution of a thermodynamic target system, e.g. during approach to equilibrium. BSM’s models are an idealisation in that they replace the actual target system with a fictitious system that can only instantiate a subset of the macro states the target system can assume. Their representation of the time evolution of the target system, in other words, is deficient in virtue of not representing all aspects of it. This idealisation entails also an approximation: BSM’s models ignore the actual macro values of the target

system's macro variables at any given time and approximate them by assigning fixed macro values for finite intervals of time.

Now this *approximation* is unproblematic. (I figure the presumption that DP might not actually be a problem originates from this fact.) If the model is too crude, if the approximation is too inexact, we may always impose a more fine-grained partition. However, that BSM's models are idealisations with respect to the possible macro states a target system can take on becomes problematic when adopting a realist stance towards BSM and taking it to be the fundamental theory, as Frigg & Werndl do. Working "under the assumption that the world is governed by Newton's equation of motion", they regard BSM as a "true complete fundamental theory". As Norton points out, one would expect the less fundamental theory to be an idealisation of the fundamental theory, not the other way around (Norton 2012, 211): "Realists, however, regard the antecedent theory as an idealization of the successor theory."

This is not to argue against Frigg & Werndl's position. I think it is well justified to regard BSM as the more fundamental theory. And if I had to be realist about either standard BSM or classical, i.e. pre-statistical, thermodynamics, I'd choose the former over the latter. But that doesn't do away with the fact that there is something odd with BSM, that its models are idealisations with respect to the possible macro states a system can take on. This oddness is the discontinuity problem of the standard framework of BSM. I think it can be avoided by adopting a new framework, as presented below, and those who want to be realists about BSM in particular should be eager to consider it.

III.5 Hints of digression from the usual

As we have seen, the construal of phase space partitioning within the framework of BSM, with sharp boundaries between a finite number of non-overlapping macro regions of finite, non-zero volume, leads to the discontinuity problem. At the same

time, this construal seems to be standard. However, there are hints to be found in the literature, showing that not everyone considers it to be entirely unproblematic. Goldstein et al. (2019, 4) remark that phase space partitioning is not as straightforward a procedure as one might think: " $\Gamma(X)$ is the set of all phase points that 'look macroscopically the same' as X . Obviously, there is no unique precise definition for 'looking macroscopically the same,' so we have a certain freedom to make a reasonable choice". And later (ibid., 28), introducing the idea of "'fuzzy' macro sets", i.e. fuzzy macro regions: "The point here is to get rid of the sharp boundaries between [macro regions] as the boundaries are artificial and somewhat arbitrary anyway." I agree with Goldstein et al.'s verdict that boundaries are "artificial". As we have seen above, they are an artefact, resulting from the way the framework of BSM is set up.

Here is Wallace (2018, 19) taking a similar line regarding the arbitrariness of phase space partitioning in BSM: "the macrostate partition at the heart of Boltzmannian statistical mechanics is [...] vulnerable to these criticisms [(concerning the adequacy of coarse-graining), A.E.] [...] Consider some standard descriptions of the coarse-graining: '[W]e must partition [phase space] into compartments such that all of the microstates X in a compartment are macroscopically indistinguishable[.]' (Callender 1999, p.355). 'Everyday macroscopic human language (that is) carves the phase space of the universe up into chunks.' (Albert 2000, p.47) If pushed, I suspect Boltzmannians would reply that it is not the epistemic indistinguishability of macrostates that is doing the work, but rather the possibility of writing down robust higher-level dynamics in terms of macrostates, and largely abstracting over microscopic details". Wallace's argument here is a tu quoque: If the Gibbsian approach can rightly be criticized for coarse-graining, then the Boltzmannian approach can rightly be criticized for partitioning phase space into macro regions. We are not concerned with the Gibbsian approach, so this discussion shall not further distract us. The point to be made is merely that others see a potential for

criticism of the BSM framework as well, and locate it at the same juncture as I do: the partition of phase space into macro regions with sharp boundaries.⁴⁸

III.6 Against sharp boundaries of macro regions in phase space (I): Crossing boundaries

Apart from a few hints and remarks like those just presented, there is, at least to my knowledge, a surprising lack of literature discussing the issue at hand. In particular, I haven't found the discontinuity problem being made explicit anywhere. In the following, I shall therefore try to make it apprehensible myself, through a series of examples revealing its consequences.

Imagine a system of interest, e.g. a gas in a (sealed, insulated) room. Modelling it in the usual way, we have a partition of the relevant, i.e. energetically available phase space Γ_E , including sharp boundaries between macro regions. These macro regions correspond to different macro states of the gas, such that, if the micro state of the gas is located within a certain macro region, the gas instantiates the macro state as

⁴⁸ Consider also the following exchange between Sheldon Goldstein and Daniel Sudarsky during the former's talk *Gibbs vs. Boltzmann Entropy* at the *Chimera of Entropy* summer school in Split in 2018 (my transcript):

Daniel Sudarsky: "Something that has bothered me all the time: consider that partition that you have [the usual partition with sharp boundaries]. Then consider a collection of points that are approaching the boundary from the two sides. I cannot imagine being able always to differentiate between them [... inaudible]"

Sheldon Goldstein: "You are right. So maybe you need to develop a better scheme which is more realistic, so you don't have sharp boundaries. [...] You are quite right; in a better world than ours, somebody will have done that. Maybe somebody has some kind of fuzzy notion of macro state. But you don't want to introduce that here [in the context of a heuristic discussion of entropy and approach to equilibrium], which will just complicate things. But you are completely right, there isn't that realistic division between [macro regions]: 'Oh, these points over here look like equilibrium, go epsilon further, doesn't look like equilibrium anymore.' That, obviously, is unrealistic. But if you want to get an understanding of the phenomena you should make that idealisation."

(<https://www.youtube.com/watch?v=jU838SXBv4>, 00:52:30 – 00:53:40)

Of course, I am aware that such a statement, made in the context of a presentation, must not be mistaken for the "official" position of the speaker. That is, I don't assume that Goldstein (or Sudarsky) would have made the same or a similar statement in a published paper. (Neither do I assume that they wouldn't make such statements in a published paper. After all, Goldstein et al. 2019 call a partition with sharp boundaries "artificial".) Nevertheless, I take Sudarsky's question and Goldstein's reply as a further hint to the effect that the way phase space is partitioned in the framework of BSM isn't unanimously viewed as entirely unproblematic.

specified by its corresponding macro values. Hence, the actual micro state of the gas determines its actual macro state. Let's say that, in its initial macro state M_0 at t_0 , the gas is confined to a certain corner of the room. Let the system evolve until the gas has spread out and evenly fills the whole room. This state we call M_{eq} and mark with the time stamp t_{eq} . Clearly, the micro state representing the system must have evolved and crossed some boundaries. To simplify the picture, imagine that there are only two macro states: M_0 with the gas in the corner, and M_{eq} with the gas spread out. Accordingly, there are two macro regions, Γ_{M_0} and $\Gamma_{M_{eq}}$.⁴⁹ Since the gas has changed its macro state from residing in the corner to filling the whole room, its micro state must have wandered from Γ_{M_0} to $\Gamma_{M_{eq}}$. Hence, it must have crossed the boundary between these macro regions. Now zoom in on the phase space region where the micro state crosses the boundary between Γ_{M_0} and $\Gamma_{M_{eq}}$. There is a precise instant when this happens, namely when the micro state lies exactly on the boundary. Let's mark this instant $t_{crossing}$. Here comes the problem: Just before $t_{crossing}$, the micro state is in Γ_{M_0} , such that it is instantiating the macro state M_0 , with the gas in the corner. Just after $t_{crossing}$, the micro state is in $\Gamma_{M_{eq}}$, such that it is instantiating the macro state M_{eq} , with the gas spread out. Hence, the continuous evolution of the micro state over just a tiny distance in phase space, such that it crosses the boundary, amounts to a rather dramatic discontinuity in the macro state of the gas: at $t_{crossing}$, the volume jumps from a rather small value (the gas confined in the corner of the room) to a comparatively large value (the gas spread out over the entire room). Yet, judging from the viewpoint of thermodynamics, as well as from considerations about the speeds at which gas particles usually move – at finite speeds, certainly – we'd expect a small, continuous evolution of the micro state only to amount to an accordingly small, and likewise continuous evolution of the macro state. After all, the continuous change of the micro state around $t_{crossing}$ might represent nothing but the displacement of one measly particle by a few microns. Such a tiny, continuous difference between micro

⁴⁹ This, of course, is an overly simplified scenario, but it is helpful to make my point.

states before and after $t_{crossing}$, one should think, can hardly account for the tremendous discontinuity between macro states M_0 and M_{eq} .⁵⁰

One might want to reply that this strange correlation between the microscopic and macroscopic evolution of the system is due to the overly simplified example we used. Certainly, there are not only two macro states, and, accordingly, not only two macro regions with one boundary between them. Our graining is way too coarse. But bear in mind: irrespective of how many macro regions we introduce in accordance with equally many macro states we care to distinguish, the boundaries between macro regions are always sharp, and macro regions are of finite number, corresponding to a finite number of different macro states of the gas. So whatever we do, as long as we have finitely many macro regions of non-zero volume, corresponding to finitely many different macro states, we will always encounter the same problem: a small, continuous change in the micro state of the gas, while crossing the boundary, will amount to a discontinuous change in its macro state.

III.7 Against sharp boundaries of macro regions in phase space (II): Shifting boundaries

Here is another, related example making the discontinuity problem explicit. This time, we are not dealing with a micro state wandering through phase space, thereby crossing boundaries between macro regions, leading to sudden, discontinuous change of the macro state. Instead, we are exploiting the arbitrariness of the partition itself. (Recall Goldstein et al.'s (2019, 4) remark above: "there is no unique precise definition for 'looking macroscopically the same,' so we have a certain freedom to make a reasonable choice".)

⁵⁰ One might have the further question: in which macro region does the micro state reside at $t_{crossing}$? Which macro state does it instantiate? In general: to which macro region do boundaries between macro regions belong? The considerations above, discussing the discontinuity at $t_{crossing}$, might suggest that the volume macro variable at this point can take on any value whatsoever, and that thus the boundary doesn't belong to any particular macro region exclusively.

Imagine a partition of the available phase space of a system, and different micro states therein. Particularly, imagine two macro regions Γ_{M_i} and Γ_{M_j} , corresponding to different macro states M_i and M_j , and separated by a sharp boundary. Now add to this picture three arbitrarily yet purposefully selected micro states, x_1, x_2, x_3 . Let's say that x_1 lies somewhere in the middle of Γ_{M_i} . Put x_2 arbitrarily close to the boundary, but still in Γ_{M_i} . x_3 shall reside equally close to the boundary, opposite from x_2 , in Γ_{M_j} . Figure 4 provides an illustration:

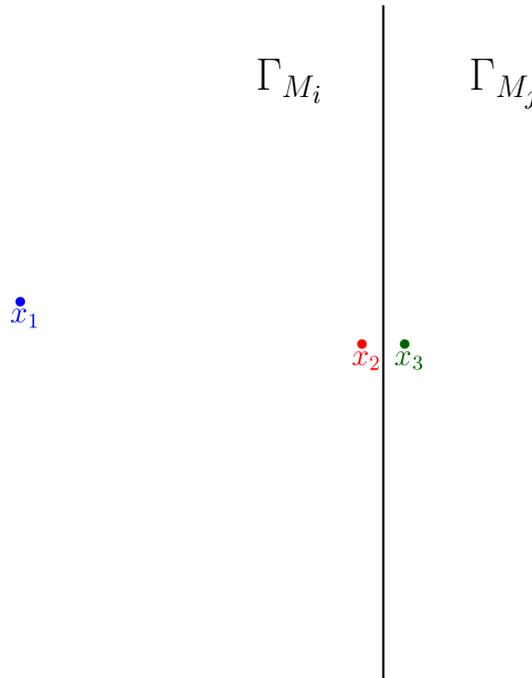


Figure 4

Just like a tiny, continuous change of the micro state shouldn't amount to anything but a tiny, continuous change of the macro state, a tiny change of the partition itself, i.e. a small shift of the boundaries of macro regions, should not amount to a model representing a completely different physical situation. Yet, in our example, this is precisely what will happen. Move, say, the boundary between Γ_{M_i} and Γ_{M_j} a wee bit towards Γ_{M_i} , such that the volume of Γ_{M_i} decreases ever so slightly, while the volume of Γ_{M_j} increases by the same amount. Make it so that now x_3 and x_2 lie in Γ_{M_j} (Figure 5).

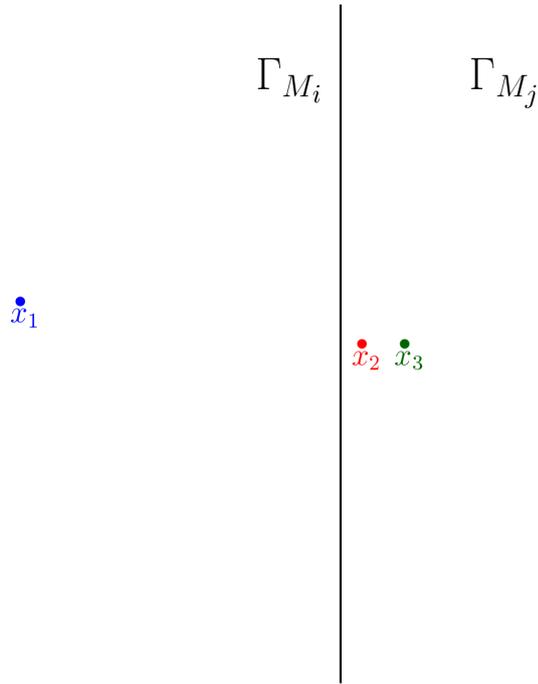


Figure 5

This new, slightly different partition seems an appropriate correction to the previous one, at least at first sight: the micro states x_2 and x_3 , always residing extremely close together, are now grouped into the same macro region Γ_{M_j} , thereby instantiating the same macro state M_j . The relatively remote micro state x_1 remains in its original region Γ_{M_i} . From a macroscopic point of view, x_2 and x_3 are now “macro-identified” in that they are taken to instantiate the same macro state. This macro state is different from the macro state instantiated by the fairly remote x_1 .

But all is not well. A tiny change of the partition, these examples show, amounts to models representing completely different physical situations on the macro level while nothing on the micro level changes: before the boundary shift, x_1 and x_2 were regarded as instantiating the same macro state M_i , while x_3 was instantiating M_j . After the shift, x_2 and x_3 are instantiating the same macro state M_j , while x_1 is still instantiating M_i . So in case x_2 becomes actual, it makes a huge difference in the model whether the boundary has been shifted or not. At the same time, nothing in the target system’s micro state has actually changed: x_2 specifies the exact same

micro state as before the shift of the boundary. The disagreement between the two models, despite describing the same microscopic situation of the same target system, is evident.

And there is more: x_1 , x_2 and x_3 have been arbitrarily selected. They are just phase points, like any other. So why not select another micro state, call it x_4 , also arbitrarily close to the (now shifted) boundary, opposite of x_2 (figure 6)?

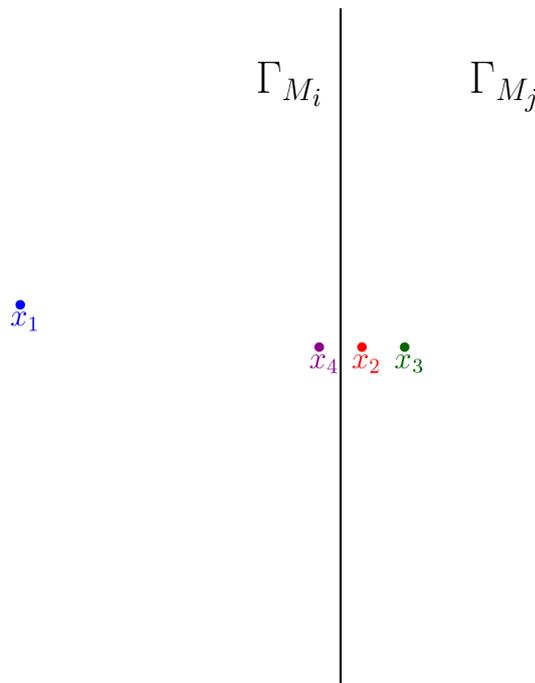


Figure 6

We face the same problem as before. So the “correction” of the partition turns out to be no correction after all. Shifting boundaries around, in effect, is just shifting around a discontinuity; it is shifting the problem to a slightly different location, thereby making the underlying problem more evident.

Sharp boundaries between macro regions in phase space, as these examples should have illustrated, are just as artificial as the border between the Netherlands and

Belgium.⁵¹ There is a line, drawn arbitrarily, but in its immediate vicinity, on both sides, things look quite the same.



III.8 Against sharp boundaries of macro regions in phase space (III): Entropy

There is yet another, related issue arising from sharp boundaries that at least deserves mentioning. Recall from above (Frigg & Werndl 2012, 102): “in fact Γ_E is almost entirely taken up by equilibrium microstates. For this reason the equilibrium state has maximum entropy.” In the usual framework, the Boltzmann entropy function is generally defined as a function assigning an entropy value to macro states, depending on the volume of their respective macro regions: $S_B(M_i) := k_B \log [\mu(\Gamma_{M_i})]$. The larger the volume, the higher the entropy. Thus, the largest region, i.e. the equilibrium region, is assigned the highest entropy. Accordingly, increase of Boltzmann entropy of a system⁵² is described by the micro state moving

⁵¹ Incidentally, the border between Belgium and the Netherlands has shifted as recently as 2018. See Stone (2018).

⁵² Boltzmann entropy of a system is defined as follows by Werndl & Frigg (2015b, 1225): “The Boltzmann entropy of a system at time t , $S_B(t)$, is the entropy of the system’s macrostate at t : $S_B(t) := S_B(M_{x(t)})$, where $x(t)$ is the system’s microstate at t and $M_{x(t)}$ is the macrostate supervening on $x(t)$.” Cf. also the almost identical definition in Frigg & Werndl (2012, 102).

to regions of larger volume (Goldstein et al. 2019, 19): “increase of Boltzmann entropy means that the phase point $[x]$ moves to bigger and bigger macro sets $[M_i]$.”

In a partition, different macro regions usually have different volumes, such that the micro state, upon traversing the boundaries between macro regions of different volume, instantiates a macro state of different entropy. So, like other macro values, as discussed above, entropy changes in a jumpy fashion. Upon the micro state moving from a smaller macro region to a larger one, entropy jumps from a lower to a higher value: the entropy function becomes discontinuous. By contrast, it stays at a constant value while the micro state wanders through a macro region, not crossing any boundaries. This behaviour, like the discontinuous behaviour of other macro values in BSM’s models, might come as a surprise. Wouldn’t we expect there to be a continuum of allowable entropy values, and that, while moving from one value to another, the continuum between them is instantiated, just like we would expect other macro values to evolve continuously? At least, it doesn’t seem implausible to say that entropy and its change *ideally* would be described by a continuous function, and not by a function that assigns the same entropy value to all micro states within one finite time interval of evolution, i.e. the time interval during which the micro state traverses a macro region, and then another, different value during the finite time interval of traversal of the micro state through another macro region, with discontinuous jumps at the boundaries. But this is how entropy behaves in the models as they result from the standard framework. Arguably, this way of modelling entropy change is at odds with the continuous macro state evolution assumed for target systems classical BSM seeks to describe. As before, the origin of the issue can be traced back to the method of partitioning phase space into a finite set of disjoint macro regions with sharp boundaries and non-zero volume.

III.9 Towards an alternative

As we have seen, phase space partitioning with sharp boundaries poses a problem for modelling continuous macroscopic change in the usual framework of BSM. I now shall revisit the examples given above while not assuming *exhaustive-cover* or *no-overlap*. As has been argued, these conditions, together with the requirement that macro regions are of non-zero volume, result in sharp boundaries. Let's see if, upon excising them, one at a time, the discontinuity problem vanishes.

III.9.1 Removing exhaustive-cover

I'd like to start with removing *exhaustive-cover*. Recall this condition: every micro state must belong to some macro region: $\forall x \exists \Gamma_{M_i} (x \in \Gamma_{M_i})$. Removing it allows for models in which there are micro states that do not belong to a macro region: $\exists x \forall \Gamma_{M_i} (x \notin \Gamma_{M_i})$. Regarding the partition, this means that there are regions of phase space that are neither a macro region, nor belong to one. Figure 7 provides an illustration:

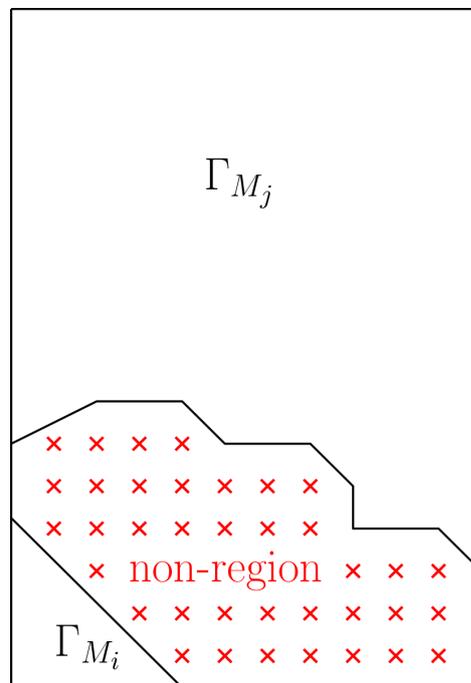


Figure 7

Via the correspondence between macro regions and macro states, phase points in those regions do not instantiate a macro state. Admittedly, a rather strange result. While wandering through phase space, it could well be under these conditions that a micro state moves from a macro region to such a “non-region” and on to a macro region again, which would mean that, in between instantiating two macro states, for some time, no macro state would be instantiated. But clearly, the constituents of the gas haven’t lost their properties, all of a sudden. The phase points within such a non-region still describe the same microscopic features of the same system as the phase points in normal regions do: all the positions and momenta of all the system’s constituents. Clearly, it is not the micro state of the gas that is responsible for this strangeness.

One might be tempted to think of these non-regions as transition regions between others. Think, again, of the gas in a room. The coarseness of the initial partitioning with only two macro regions can be reduced by introducing a “non-region” between these two. Let’s call the resulting macro regions Γ_{M_i} , Γ_{M_j} and $\Gamma_{M_{between}}$ (figure 8). That is all fine and well. The initial partition, after all, was too coarse. But observe that now, the so-called “non-region” shows its true face: it is by no means a non-region, corresponding to no macro state at all. Instead, the points it contains instantiate a macro state $M_{between}$ between the initial state M_i and the final state M_j . All the strangeness vanished. Even so-called non-regions are regions. In other words: there are no non-regions in phase space. Taking this route, *exhaustive-cover* isn’t removed after all.

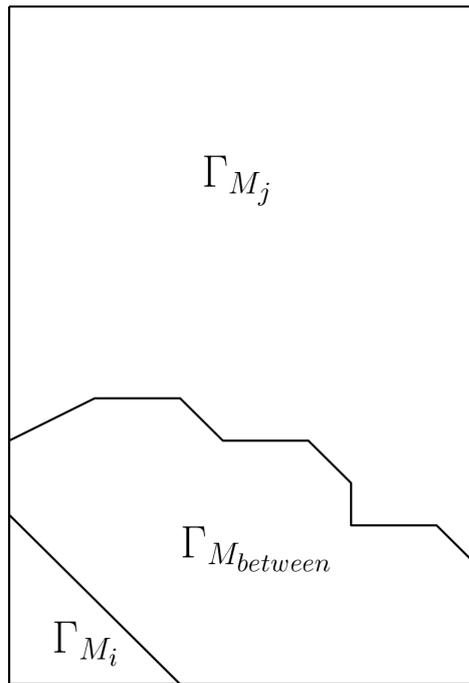


Figure 8

Apart from the strangeness of having non-regions in phase space, even if one regards them as a perfectly normal feature of BSM's models, one would be facing the same discontinuity problem as before, and even more drastically. In transitions from ordinary regions to ordinary regions, a jump discontinuity between macro values appears. In transitions from ordinary regions to non-regions and vice versa, an essential discontinuity appears: one of the limits L^- or L^+ (see above) doesn't exist, since non-regions don't correspond to macro values. So there is not just a discontinuous jump between different macro values, but a jump from a macro value to an undefined value or vice versa, making the issue explicit even more drastically.

As an interim result, it can be noted that it seems impossible to do away with the *exhaustive-cover* condition. For one, the resulting "non-regions" turned out to be either absurdities, rendering the discontinuity problem even worse, or to be ordinary regions after all. In any case, the problem we wanted to solve remains.

III.9.2 Removing no-overlap

So let's try keeping *exhaustive-cover* and excising *no-overlap* instead. *No-overlap*, recall, says that every micro state can belong to only one macro region:

$\forall x \forall \Gamma_{M_i, j} \neg (x \in \Gamma_{M_i} \wedge x \in \Gamma_{M_j}); i \neq j$. Removing this condition, hence, allows for microstates that belong to more than one macro region: $\exists x \exists \Gamma_{M_i} \exists \Gamma_{M_j} (x \in \Gamma_{M_i} \wedge x \in \Gamma_{M_j}); i \neq j$. This results, accordingly, in overlap regions of phase space that belong to two macro regions (figure 9).

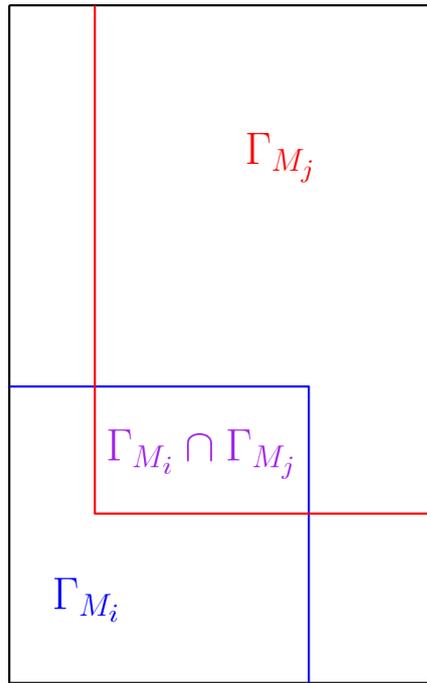


Figure 9

A micro state within such an overlap region $\Gamma_{M_i} \cap \Gamma_{M_j}$ is in Γ_{M_i} as well as Γ_{M_j} , and therefore instantiates both, M_i and M_j . As long as these have the same macro values, all is fine. However, they can't. As per usual, since macro states are specified via macro values, and correspond to macro regions, macro states being specified by the same macro values are identical macro states, i.e. there is only one macro state, corresponding to one macro region. In other words: macro regions corresponding to identical macro states form a union into one macro region. So, whenever we want to speak of different macro regions, they must correspond to different macro states. And that, in turn, means that an overlap region of two different macro

regions corresponds to two different macro states. Accordingly, every micro state in an overlap region instantiates two different macro states at the same time, which is an inconsistency.

To provide an extreme example of this equally strange result, recall our simple gas example. Imagine that there is some overlap $\Gamma_{overlap} = \Gamma_{M_0} \cap \Gamma_{M_{eq}}$. If the micro state is in $\Gamma_{overlap}$, it is in Γ_{M_0} and $\Gamma_{M_{eq}}$. Accordingly, the gas is concentrated in the corner (M_0) as well as evenly spread out over the whole room (M_{eq}) – an inconsistency even worse than the problem we are trying to solve. Hence, one should abstain from excising *no-overlap* as well.

III.9.3 A different solution?

Removing either *exhaustive-cover* or *no-overlap* from the usual framework results in severe problems. So they should be kept. Yet, at the same time, leaving everything at status quo provides no solution to the discontinuity problem as it arises from sharp boundaries between a finite number of macro regions of non-zero volume. Consequently, we must ask ourselves whether a framework in which a set of phase points, i.e. a macro region, corresponds to *one and the same* macro state is a good one. Tweaking the framework's conditions didn't result in the solution we had hoped for. Maybe, then, it is time to consider adopting a different one, which is what I am proposing in the following section.

III.10 From macro regions to worlds

Instead of disjoint macro regions as described above, I suggest employing an alternative way of carving up phase space – something that, to some extent, resembles macro regions, yet differs in important respects. In order to avoid the risk of confusion with ordinary macro regions, these new sets of phase space points shall be called *worlds*. They will be defined by a relation of *sufficient similarity* (SS), in a way to be made precise shortly. As a result, there no longer will be a partition

into a finite number of macro regions of non-zero volume and with sharp boundaries, where different macro regions stand in a one-to-one correspondence to different macro states. Once the idea that we don't need such a partition has grown upon us, an adequate correspondence between the continuous evolution of the micro state and the thermodynamically expected, continuous evolution of the macro state can be recovered.

Contrary to the micro states contained in a standard macro region, the micro states in a world do not all instantiate the *same*, i.e. *identical* macro states. (Some do, namely those representing permutations of identical constituents.)⁵³ Instead, the construal of worlds is based on the notion of *indistinguishable* macro states, which defines a broader class: identical macro states are certainly indistinguishable, but indistinguishable macro states are not necessarily identical. This, to some, might seem as an unfavourable introduction of subjectivity⁵⁴ into the framework.

Speaking of indistinguishable macro states triggers the question: indistinguishable for whom or what? However, replacing sameness with indistinguishability allows the resulting models to represent continuous macroscopic change, in accordance with the continuous evolution of the micro state, and with target systems for which a thermodynamic description is adequate. So there clearly are some advantages. In the end, it will be for the reader to decide whether they prefer a framework that contains no reference to distinguishability, but is unable to model continuous macroscopic change, or a framework that is able to model continuous macroscopic change, but refers to indistinguishability.

Just how exactly this new framework is built and supposed to work is the topic of this section.

⁵³ For a discussion of the permutation of identical constituents/particles, see below.

⁵⁴ The term "subjectivity" might actually be a bit misleading. What I am going to provide is an objective description that takes into account matters of perspective. There is nothing subjective about this description itself. Cf. Hemmo & Shenker (2012), who emphasize that macro states, and hence partitioning into macro regions, are objective, despite being observer-relative, and Callender (1999, 370), pointing out that "the coarse graining used to define [Boltzmann entropy] does not imply that [Boltzmann entropy] is subjective, only that it is relational".

III.10.1 Introducing a new framework for BSM: Some conditions

Most importantly, by developing this new framework, I want to avoid the discontinuity problem as it arises from the old framework. To achieve this goal, my construal of worlds as sets of micro states in phase space should avoid the imposition of sharp boundaries between different worlds, each corresponding to a particular macro state, such that macro values jump discontinuously upon the micro state crossing such boundaries. That, however, is not to say that worlds may not have *any* boundaries at all. In fact, they may even have sharp boundaries. But these boundaries must differ from those we encounter in the old framework in that crossing them does not result in discontinuous macroscopic change. This will be made possible by allowing worlds to overlap, without micro states in the overlap instantiating inconsistent macro states, as is the case for overlapping macro regions in the old framework (see section *Removing no-overlap*). In order to avoid such inconsistencies, a micro state in an overlap must not instantiate *different* macro states. As shown above, micro states in overlapping macro regions would instantiate different macro states (which is why there can't be overlapping macro regions in the old framework). This is due to the fact that the finite set of different macro regions is bijectively mapped onto the finite set of different macro states: each macro region corresponds to exactly one macro state, and vice versa. So, to avoid inconsistent instantiations in the overlap of worlds, worlds must not correspond to exactly one macro state, and macro states must not correspond to exactly one world. Rather, worlds shall correspond to a set of macro states that are indistinguishable⁵⁵ from a given macro state. This set can be a continuum of macro states, which, as we will see, enables modelling continuous macroscopic change. Of course, the reasonable stipulation that every micro state instantiates exactly one macro state will not be touched, for this would undermine supervenience of the macro state on the micro state. Accordingly, since worlds are sets of micro states,

⁵⁵ Relative to a given "resolution power", see below.

and are supposed to correspond not to one, but a continuum of macro states, the micro states within a world cannot all instantiate the same macro state, as the micro states in macro regions do. Instead, micro states within a world will be related in that they all instantiate macro states that are indistinguishable from a given macro state. This feature will be expressed in the relation “sufficient similarity” (SS), pertaining to micro states: a world, then, will be defined as a set of sufficiently similar micro states, such that all micro states in a world instantiate a macro state that is indistinguishable from a given macro state.

It should be mentioned that already in the old framework, macro states trivially can be indistinguishable – namely if they are identical. Because of that, and because of how the relation SS is defined via indistinguishability, micro states in the old framework can be sufficiently similar as well. In fact, since all micro states in a macro region instantiate the same macro state, all micro states in a macro region are sufficiently similar to all other micro states in the same macro region. Here, we already encounter an important difference between worlds and macro regions. The micro states in a world are not necessarily sufficiently similar to all other micro states in that world. It is merely required that they be sufficiently similar to what I call the “central” micro state of a world, around which the world is defined. Take, for example, some arbitrary micro state. It certainly instantiates a macro state. If we want to construe a world around this micro state, we merely require that all micro states in this world be sufficiently similar to that micro state in that they instantiate a macro state that is indistinguishable from the macro state instantiated by the selected micro state. This does not entail the requirement that all micro states in the thusly construed world instantiate macro states indistinguishable from the macro states instantiated by all other micro states in that world. It is very much allowable that some micro states in a world instantiate macro states that aren’t mutually indistinguishable. Hence, not all micro states in a world are necessarily sufficiently similar to all other micro states in the same world. At the same time, micro states from different worlds may be sufficiently similar in that they

instantiate indistinguishable macro states. In other words: both relations, indistinguishability as well as sufficient similarity, are nontransitive.

Let's make things reasonably precise.

First, supervenience of the macro state on the micro state is claimed. That is, if the macro state changes, the micro state must change as well. In other words: every micro state instantiates exactly one macro state.⁵⁶ Formally:

$$\forall x_i \exists! M_i: (M_i x_i) \quad \text{(supervenience)} \quad ^{57}$$

Read: For every micro state x_i there is *exactly one* macro state M_i such that x_i instantiates M_i .

Like before, macro states are defined via sets of macro values. Saying that macro states are indistinguishable means that it is impossible to tell them apart based on these macro values. So the question arises: when, exactly, is it impossible to tell macro states apart based on their macro values (and when is it possible)? When are macro values similar enough? In order to answer these questions, I must elaborate the way in which macroscopic distinguishability depends on “resolution power”⁵⁸.

Saying that macroscopic (in)distinguishability depends on resolution power means the following: observers and measuring instruments have certain limitations with respect to their capabilities to differentiate between different macro states.

⁵⁶ According to Frigg & Werndl (2011, 125), supervenience of the macro state on the micro state is “a basic posit of the Boltzmann approach”. In this respect, the new framework doesn't differ from the old one. See also Frigg (2008, 3.2.1) and Werndl & Frigg (2015a, passim), and recall Frigg et al. (2016, 4.1) from above: “to every given microstate x there corresponds *exactly one* macrostate”.

⁵⁷ $\exists!$ denoting unique existential quantification – “there is exactly one” – as opposed to normal existential quantification with the quantifier \exists – “there is (at least) one” – allowing for the possibility that there is more than one. Unique existential quantification is used for brevity. It can be expressed in terms of normal existential quantification: instead of formulating the supervenience condition as $\forall x_i \exists! M_i: (M_i x_i)$, one can equivalently write $\forall x_i \exists M_i: (M_i x_i \wedge \neg \exists M_j (M_j x_i \wedge M_i \neq M_j))$.

⁵⁸ Borrowing the term from Hemmo & Shenker (2012), without necessarily meaning the exact same concept as they do.

Depending on their abilities/precision, to which I refer as their resolution power, two given macro states, specified by different macro values, might be distinguishable for a given observer/instrument, or not. Thus, macroscopic (in)distinguishability is observer- or instrument-relative in that it depends on the abilities or limitations of the observer or instrument: relative_{observer}⁵⁹ to a given observer and/or instrument with a certain resolution power, certain macro states are distinguishable (or not). Accordingly, for indistinguishable macro states, it is not claimed that they are specified by identical macro values, as it is done for same macro states in the old framework. Indistinguishable macro states may differ in their macro values. For example, the macro states M_1 and M_2 might differ in temperature values $T_1 = 23.542^\circ\text{C}$ and $T_2 = 23.543^\circ\text{C}$. But with a limited resolution power on the order of $\Delta T = 0.1^\circ\text{C}$, they are, with respect to temperature, rendered indistinguishable, despite being different. Already on a conceptual level, it seems reasonable to not regard *different* macro states as the *same* macro state, even if they are indistinguishable. Macroscopic indistinguishability being relative_{observer} comes on top: in principle, a being like Laplace's Demon can distinguish between all possible micro states, and thus all macro states. And in the example given, even a human observer can render the macro states in question distinguishable, if they increase their resolution power, for example by using a better thermometer. Thus, it seems unreasonable to identify macro states based on their indistinguishability, which is a contingent, observer-relative, matter of fact. Nevertheless, macro states can be grouped together via their indistinguishability. For example, if one measures or prepares the temperature of a system to be 23.5°C in accordance with a certain resolution power, one may pretend that this is the precise macro value, namely $23.5000\dots^\circ\text{C}$. Thus, a precise macro state is specified, even though it is practically impossible to measure or control to absolute precision. Generally speaking, the *formal* specification of macro states via macro values is not limited by resolution power: one is free to be arbitrarily precise when specifying macro states. Irrespective of that, due to limited

⁵⁹ I write "observer" as a shorthand for "observer and/or instrument".

resolution power, other macro states, differing in macro values only to an extent that cannot be resolved, are rendered indistinguishable from the precise macro state specified. Via this indistinguishability, macro states can be grouped together with a given macro state, without identifying them with the latter, or with one another.

Having established that macroscopic (in)distinguishability is observer-relative, the question above can be answered: macro states are distinguishable iff, given a certain resolution power, an observer/instrument can distinguish between their respective macro values. For example, M_1 and M_2 might actually have different temperature macro values $T_1 = 23.542^\circ\text{C}$ and $T_2 = 23.543^\circ\text{C}$, but given a resolution power of $\Delta T = 0.1^\circ\text{C}$, these macro values can't be distinguished through observation or measurement. In this sense, they are indistinguishable.

With the notion of macroscopic (in)distinguishability made clear, worlds can now be introduced into the framework. First, sufficient similarity between micro states is defined via macroscopic indistinguishability of macro states (where “ \approx ” denotes macroscopic indistinguishability):

$$\forall x_{i,j}: \left(x_j \text{SS} x_i \leftrightarrow \exists M_{i,j} \left(M_j x_j \wedge M_i x_i \wedge (M_j \approx M_i) \right) \right) \quad (\text{sufficient similarity})$$

Read: For all pairs of micro states $x_{i,j}$: x_j and x_i are sufficiently similar iff there are macro states M_i and M_j such that M_j is instantiated by x_j , M_i is instantiated by x_i , and M_j and M_i are indistinguishable.

It is then possible to construe a world \mathcal{W}_i as the set of all x_j that satisfy SS with respect to a given x_i .⁶⁰

⁶⁰ One might think that defining worlds *around* a central micro state is somewhat odd. But the idea is not at all new. Take Reichenbach (1956, 71) for an example, even calling this method “natural”: “The distinction between microstates and macrostates [...] leads to a natural grouping of microstates. Given a microstate a_1 , the macrostate A to which it belongs is well determined: A includes all those microstates a_i which are macroscopically indistinguishable from a_1 .” Note the slight differences in

$$\forall x_i \exists \mathcal{W}_i \left(\forall x_j (x_j \in \mathcal{W}_i \leftrightarrow x_j SS x_i) \right) \quad (\text{world membership})$$

Spelled out in terms of indistinguishable macro states, *world membership* reads:

$$\forall x_i \exists \mathcal{W}_i \left(\forall x_j \left(x_j \in \mathcal{W}_i \leftrightarrow \exists M_{i,j} \left(M_j x_j \wedge M_i x_i \wedge (M_j \approx M_i) \right) \right) \right)$$

For all micro states x_i there is a world \mathcal{W}_i such that, for all micro states x_j , x_j is a member of \mathcal{W}_i iff there are macro states M_i and M_j such that M_j is instantiated by x_j , M_i is instantiated by x_i , and M_j and M_i are indistinguishable.

Under this construal, micro states instantiating indistinguishable macro states belong to the same world.

A few remarks are in order before moving on. First, it is worth stressing that SS is a reflexive and symmetric but nontransitive relation. It inherits these relational properties from the notion of macroscopic indistinguishability, itself being nontransitive. Hence, worlds, as defined via SS, are not equivalence classes, as disjoint macro regions in a partition are. While all micro states x_j in a world \mathcal{W}_i are sufficiently similar to the central micro state x_i , it is not necessary that all x_j in said world *mutually* satisfy SS.⁶¹

terminology. What Reichenbach calls a “macrostate” is a set of micro states and thus would be called a *macro region* in the old framework. But actually, since it is only required that the a_i be macroscopically indistinguishable from a given a_1 , i.e. that they satisfy SS with respect to a_1 , the resulting set is not a *macro region*, but a *world* in our sense: not all a_i necessarily are macroscopically indistinguishable on Reichenbach’s construal. Unfortunately, Reichenbach didn’t seem to realize that, once worlds – macrostates in his terminology – are defined around a given micro state via a similarity relation, it is only a small step towards overlapping worlds.

⁶¹ For an illustrative example, think of the height of people: When Amy and Berta are of the same height, and Berta and Chris are of the same height, then necessarily Amy and Chris are of the same height. Sameness is a reflexive, symmetric and transitive relation. Not so similarity. If Amy and Berta are of similar height, and Berta and Chris are of similar height, then it could be that Amy and Chris are of similar height, but it is by no means necessary. Hence, other than sameness, similarity is reflexive and symmetric but nontransitive. Likewise, the case of macroscopic indistinguishability: even if I am unable to notice the difference in temperature of a gas at different times t and t' , and

Second, x_i is called the central micro state of the world \mathcal{W}_i . In this vein, every micro state is the central micro state of its own world, and conversely, every world has a central micro state. One might want to know how to find the central micro state and its world, given only a macro state, specified via certain macro values.

Practically, due to macro states being distinguishable relative_{observer} to a certain resolution power, this is impossible, but also unnecessary. One can easily specify a macro state to arbitrary precision, simply by fixing macro values. By contrast, it is not at all easy to specify a precise micro state: one would have to fix the values of all (relevant) degrees of freedom of all constituents of the system of interest in order to specify a unique point in phase space. This is unfeasible. The best one can do practically is to prepare a system such that it is in a macro state indistinguishable from the specified macro state, according to a given resolution power. The actual micro state of the system so prepared is then sufficiently similar to other micro states which, counterfactually, instantiate macro states that are indistinguishable from the specified one. Alternatively, one can measure the macro values of the system and take them as the specification of the macro state. The system will certainly be in some micro state, instantiating a macro state indistinguishable from the one specified. For all practical purposes, this is enough. If one wants to close in on a specific micro state, the preparation of the system must become better, or resolution power must be increased, such that the world of sufficiently similar micro states becomes smaller. But since practically, resolution power is limited, it is practically unattainable to find a specific micro state this way.

Let's move on with building the framework. Earlier, for the old framework, we have seen that regions of phase space not corresponding to any macro state – what I have called non-regions – are rather absurd: every point in phase space precisely

unable to notice the difference in temperature of the same gas at t' and t'' , I might be able to notice the difference in temperature of the gas at t and the gas at t'' . Hence, despite the gas at t and t' being macroscopically indistinguishable (with respect to temperature), and the gas at t' and t'' being macroscopically indistinguishable, it does not follow that the gas at t and t'' is macroscopically indistinguishable.

describes a possible micro state and thus, via *supervenience*, (counterfactually) instantiates a macro state. But as soon as a micro state instantiates a macro state, it is part of a macro region (in the old framework), or of a world (in the new framework). Hence, every possible micro state is part of a world. This is called *exhaustive cover*:

$$\forall x_i \exists \mathcal{W}_i (x_i \in \mathcal{W}_i) \quad (\text{exhaustive cover})$$

Exhaustive cover immediately follows from *world membership*, *sufficient similarity* and *supervenience*: take *world membership* and set $x_j = x_i$, such that $\forall x_i \exists \mathcal{W}_i (x_i \in \mathcal{W}_i \leftrightarrow x_i S S x_i)$. That $x_i S S x_i$ is satisfied can be seen from its analysis in *sufficient similarity*: $\forall x_i: (x_i S S x_i \leftrightarrow \exists M_i (M_i x_i \wedge (M_i \rightleftharpoons M_i)))$. That, for every x_i , there is an M_i such that it is instantiated by x_i is seen from *supervenience*, and $M_i \rightleftharpoons M_i$ is trivial. So $\forall x_i \exists M_i (M_i x_i \wedge (M_i \rightleftharpoons M_i))$ is satisfied. So x_i is sufficiently similar to itself: $x_i S S x_i$. And since the quantification ranges over all x_i , all x_i are members of the World \mathcal{W}_i around x_i , i.e. of their own world. Thus, phase space is exhaustively covered by worlds.

However, this is not to say that every micro state belongs to only one world. Via *world membership*, spelled out in terms of macroscopic indistinguishability, all micro states instantiating macro states indistinguishable from a given macro state belong to the same world. Phase space is covered by worlds, and every micro state is a member of its own world. With these worlds also containing other micro states, it immediately follows that worlds overlap. The possibility of a micro state being a member of different worlds is one of the discerning features of the new method of carving up phase space. Iff a micro state x_k belongs to worlds $\mathcal{W}_i, \mathcal{W}_j$, it lies in their overlap $\mathcal{O}_{\mathcal{W}_i, \mathcal{W}_j} = \mathcal{W}_i \cap \mathcal{W}_j$:

$$\forall x_k \forall \mathcal{W}_i, \mathcal{W}_j (x_k \in \mathcal{O}_{\mathcal{W}_i, \mathcal{W}_j} \leftrightarrow x_k \in \mathcal{W}_i \wedge x_k \in \mathcal{W}_j) \quad (\text{overlap})$$

Like *world membership* before, *overlap* is spelled out in terms of indistinguishable macro states:

$$\forall x_{i,j,k} \forall \mathcal{W}_{i,j} : \left(x_k \in \mathcal{O}_{\mathcal{W}_{i,j}} \right. \\ \left. \leftrightarrow \exists M_{i,j,k} (M_k x_k \wedge M_i x_i \wedge M_j x_j \wedge (M_i \approx M_k) \wedge (M_j \approx M_k)) \right)$$

From this, it can be easily seen that, if there are micro states x_i , x_j and x_k , instantiating – via *supervenience* – macro states M_i , M_j and M_k , and M_k is indistinguishable from both M_i and M_j , then there is an overlap of worlds \mathcal{W}_i and \mathcal{W}_j containing the micro state x_k . It is important to note that for this result, it is not necessary to assume that M_i and M_j are indistinguishable.

The converse can be seen just as easily: if there is an overlap containing an x_k , then there is a macro state M_k that is indistinguishable from both M_i and M_j , even if the latter two are distinguishable.

Figure 10 illustrates these results. Two micro states, x_i and x_j , instantiate distinguishable macro states, M_i and M_j . One can think of this picture as superimposed snapshots of the micro evolution of a system, at different times. Say, for example, that x_i is the micro state of the system at time t , and that x_j is the micro state of the system at a later time t'' . The picture represents a case in which our resolution power is good enough to distinguish between the macro states instantiated by x_i at t and x_j at t'' . Accordingly, x_i is not in \mathcal{W}_j , which would render it sufficiently similar to x_j , and likewise, x_j is not in \mathcal{W}_i . During its evolution from x_i to x_j in the time interval $[t, t'']$, the micro state traverses the overlap $\mathcal{O}_{\mathcal{W}_{i,j}}$. Let's pick out x_k as the micro state at time t' , between t and t'' . Any x_k in the overlap will do. Since it is in the overlap, it instantiates a macro state M_k , that, with the given resolution power, is indistinguishable from both, M_i and M_j .

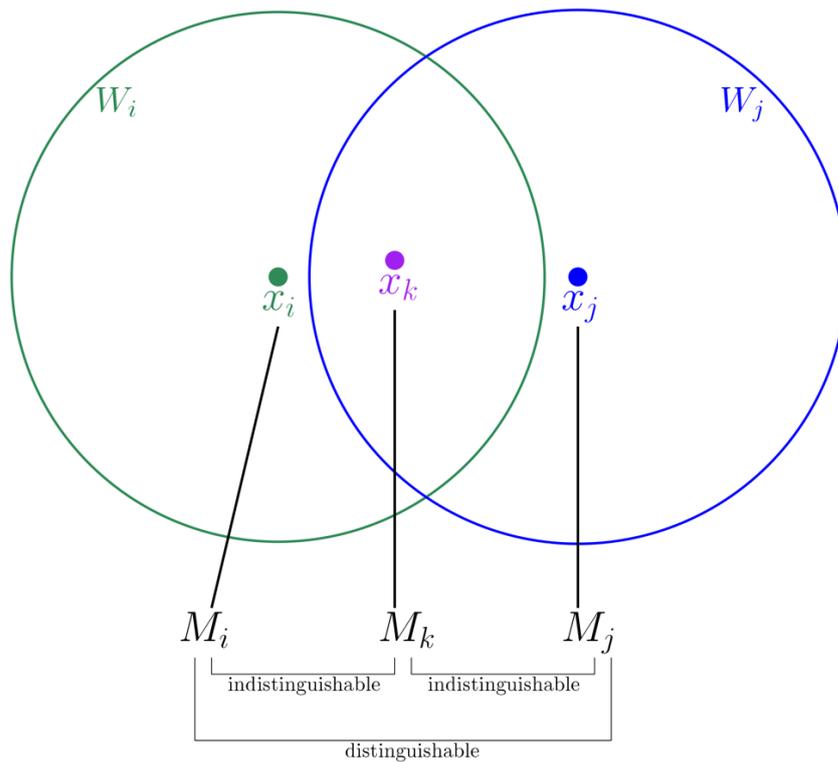


Figure 10

The phase points x_i , x_j and x_k in figure 10 are drawn at a significant distance, because they each represent the micro state of a system at different times t, t', t'' . The picture drawn is an overlay of “snapshots” at these times. But it is easy to imagine the continuous trajectory of the actual evolution of the micro state through phase space during the time interval $[t, t'']$ instead. This trajectory being continuous – as is assumed for the micro-state evolution in the old as well as in the new framework – contains an uncountable infinity of micro states. And since all these micro states instantiate a macro state, the uncountably infinite set of micro states in the trajectory bijectively corresponds to an uncountably infinite set of macro states instantiated, and thus, to an uncountably infinite set (of sets) of macro values bijectively corresponding to the set of macro states. So the issue that gave rise to the discontinuity problem in the old framework – the issue that there was an uncountable infinity of micro states that couldn’t be mapped bijectively onto a finite set of macro states, and the set of macro states had to be finite because there had to be a finite set of macro regions, and the set of macro regions had to be finite because the available phase space is finite and macro regions had to have non-zero

volume, and macro regions had to have non-zero volume because the Boltzmann entropy function assigns entropy values to macro states based on the volume of their corresponding macro regions, and entropy of different macro states is required to take on different values in order to even speak of entropy increase – this issue seems to have vanished completely. In the new framework just proposed, we likewise want to be able to speak of entropy increase, so we likewise have to have worlds of different volumes, so we likewise have to have worlds of non-zero volume. But by contrast, models in the new framework aren't restricted to a finite number of worlds by these requirements, because there are as many worlds as there are micro states, and there are as many worlds as there are micro states – despite worlds having non-zero volume – because worlds can overlap, and worlds can overlap because the micro states they contain don't all instantiate the *same* macro state, but merely ones that are indistinguishable from a given macro state. Because of all this, change of macro values doesn't have to be modelled as discontinuous in the new framework, as opposed to the old one. Using macroscopic indistinguishability instead of sameness, it is possible to avoid models in which the evolving micro state, in order to instantiate different macro states, must cross a boundary at which macro values jump. In models resulting from the new framework, the micro state can wander through phase space without crossing such boundaries, instantiating indistinguishable macro state after indistinguishable macro state, until, at the two far ends of its continuous trajectory, it does instantiate distinguishable macro states.

One might be tempted to think that accounting for limited resolution power would solve the discontinuity problem already in the old framework: simply make the partition so fine-grained that neighbouring macro regions correspond to different, yet indistinguishable macro states. Then, the transition of the micro state from one to the other wouldn't result in a noticeable discontinuity. This latter claim, speaking of *noticeable* discontinuities, is correct. However, it doesn't solve the discontinuity problem. DP doesn't arise merely because of *noticeable* discontinuities, but because of discontinuities *simpliciter*, whether noticeable or not; because the old framework

cannot possibly model continuous change of macro values. It can't because there is only ever a finite set of allowable macro values in a given model, and a finite set of macro values cannot possibly be mapped bijectively onto a continuum of macro values as present in thermodynamically described target systems, and likewise, onto a continuum of micro states in a trajectory through phase space. This remains an idealisation – no matter how fine-grained the partition – that is problematic especially when regarding BSM as a fundamental theory (see above). By contrast, the new framework is able to model continuous change of macro values because it actually allows for a continuum of macro values to exist in its models, and this continuum can be mapped bijectively onto a continuum of micro states in a trajectory, and, trivially, onto a continuum of macro values as present in target systems that are adequately described by thermodynamics.

In short, as long as the micro states in a macro region are all supposed to instantiate the same macro state, the discontinuity problem arises. The new framework doesn't solve it merely by taking into account resolution power, but by replacing macro regions with worlds, sets of micro states that instantiate macro states *indistinguishable* from a given macro state, but not necessarily the same, such that worlds can correspond to a continuum of macro values. Indistinguishability between macro states alone doesn't imply that a set of macro values is a continuum. But at least, the new framework is able to model such cases. In solving the discontinuity problem, all the crucial work is done by replacing a finite set of macro regions with an uncountable infinity of worlds.

So much for a general overview of the proposed framework. In the following section, I would like to demonstrate its capabilities by applying it to the familiar example of a gas spreading out in a room.

III.10.2 Back to the example

Suppose there is an initial micro state x_0 instantiating M_0 with the gas concentrated in the corner of the room. The set of sufficiently similar micro states, via *world membership*, forms the world \mathcal{W}_0 around x_0 . Let x evolve away from x_0 , such that the gas is spreading out: x_1 instantiates a slightly, but noticeably more spread out macro state of the gas, M_1 . Since M_1 is macroscopically distinguishable from M_0 , x_1 is not a member of \mathcal{W}_0 . But it certainly is a member of its own world, \mathcal{W}_1 . Despite x_1 and x_0 residing in different worlds \mathcal{W}_0 and \mathcal{W}_1 , we can select $x_\mathcal{O}$ such that their worlds do overlap: there are (uncountably many) micro states between x_0 and x_1 that meet the *world membership* condition of both, \mathcal{W}_0 and \mathcal{W}_1 . They make up the overlap $\mathcal{O}_{\mathcal{W}_0,1}$. Let's pick out one of them, $x_\mathcal{O}$. The macro state $M_\mathcal{O}$ it instantiates is indistinguishable from M_0 , such that $x_\mathcal{O}$ is part of \mathcal{W}_0 . But $M_\mathcal{O}$ is also indistinguishable from M_1 , such that $x_\mathcal{O}$ is part of \mathcal{W}_1 . So, when x evolves from x_0 to $x_\mathcal{O}$, and hence the instantiated macro state changes from M_0 to $M_\mathcal{O}$, this change is unnoticeable. When x then moves from $x_\mathcal{O}$ to x_1 , the macro state changes from $M_\mathcal{O}$ to M_1 , but again, this change is unnoticeable. Yet now, x instantiates M_1 , which is noticeably different from M_0 . So, through a series of unnoticeable changes, the macro state has changed noticeably. In this fashion, we can describe the gas in its evolution from residing in the corner to being spread out over the whole room, through a series of unnoticeable changes that together account for the very noticeable difference between the initial state and the equilibrium state. And remember that the micro states we are referring to in this example are merely selected snapshots from a continuous trajectory, containing an uncountable infinity of actual micro states, instantiating an uncountable infinity of macro states. From such an uncountably infinite set of micro states, we can always select a finite set of three micro states that behave in the way just described for x_0 , $x_\mathcal{O}$ and x_1 .

III.10.3 World size, the equilibrium world, and entropy increase

So much for the continuous change between macro states. But how about the sizes of worlds? In particular, is \mathcal{W}_{eq} , like $\Gamma_{M_{eq}}$, still dominant, i.e. “vastly larger than any other [world]”, such that the accessible phase space volume is “almost entirely

taken up by equilibrium microstates”? Given the special interest in M_{eq} , will x_{eq} and its world \mathcal{W}_{eq} turn out to be special? Can the approach to and remaining in this special world still be explained in this framework? Is the change of entropy over time now continuous? A few considerations are important.

First, we must state clearly what determines Boltzmann entropy. In the old framework, it was assumed that all micro states in a macro region instantiate the same macro state. Endowing phase space with a suitable measure for macro regions (Frigg & Werndl 2011, 125f.), “the Boltzmann entropy of a macro state M [...] measures the portion of the system’s γ -space [i.e. phase space] that is taken up by microstates that correspond to M ”. In other words, it depends on the volume of the macro region corresponding to a given macro state. Thus, Boltzmann entropy can be assigned directly to a macro state (Frigg & Werndl 2012, 102): “The Boltzmann entropy of a macrostate M_i is $S_B(M_i) := k_B \log[\mu(\Gamma_{M_i})]$ (where k_B is the Boltzmann constant), and the Boltzmann entropy of a *system* at time t , $S_B(t)$, is the entropy of the macrostate of the system at t : $S_B(t) := S_B(M_{x(t)})$, where $x(t)$ is the microstate at t and $M_{x(t)}$ is the macrostate corresponding to $x(t)$.”

For the new framework, we can retain the latter definition of the entropy of a system at time t , but the definition of the entropy of a macro state must be slightly altered: it must depend on the volume of *worlds*, not on the volume of macro regions. Accordingly, Boltzmann entropy in the new framework is assigned to a given macro state M_i by considering not only M_i itself, but all macro states that are indistinguishable from M_i . This means that we have to take into account the set of all micro states that instantiate macro states indistinguishable from M_i , i.e. the world \mathcal{W}_i : $S_{B(new)}(M_i) := k_B \log[\mu(\mathcal{W}_i)]$.

It is a distinctive feature of the new framework with overlapping worlds that the sum of the volume of all worlds is infinite. Since every micro state has a world of non-zero volume around it, and there are infinitely many micro states in a given, finite volume of available phase space, there are infinitely many worlds of non-zero

volume in a finite volume of phase space. But this does not impede the possibility of comparing the sizes of individual worlds with each other, or with the volume of the available phase space. For example, for a particular equilibrium macro state, we should still be able to say that the world consisting of all micro states instantiating macro states indistinguishable from it is larger than worlds consisting of micro states instantiating macro states indistinguishable from a given non-equilibrium macro state, and thus assign a higher entropy value to the former. Likewise, since the volume of worlds remains comparable to the volume of the available phase space in total, we should still be able to say that, for example, a particular equilibrium world covers almost all of the available phase space volume, while other, non-equilibrium worlds, don't.

Note that the equilibrium world \mathcal{W}_{eq} is defined around the micro state x_{eq} that instantiates the equilibrium macro state M_{eq} . There is nothing that enforces the selection of this particular equilibrium micro state, or this particular equilibrium macro state. After all, there is a continuum of them. As long as a micro state instantiates a macro state indistinguishable from equilibrium, any will do. Nevertheless, we might be able to search for the "best" x , if we want or need to. Just select one micro state that instantiates a macro state indistinguishable from equilibrium, and then shift it around, until, via *world membership*, the volume of its world is maximized. This way, a unique "best" x and a unique equilibrium world is selected. The thusly selected equilibrium micro state may be called the "true" equilibrium micro state. Of course, whether the actual micro state of a system in macroscopic equilibrium is precisely the true equilibrium micro state (with almost certainty, it isn't) or just one sufficiently similar is irrelevant for all practical purposes.

Lastly, approach to and remaining in equilibrium for long time periods have essentially the same explanation in the new framework as in the old: because of their relative sizes, the micro state is more likely to evolve from smaller worlds into larger ones, and ultimately to end up in the equilibrium world, where it tends to

stay. It goes without saying that, nevertheless, fluctuations out of equilibrium are still possible, despite being incredibly unlikely. Yet, the new framework differs in this respect, too. Remember the issue of discontinuous Boltzmann entropy increase addressed above: with the micro state moving from one macro region to another, both being differently sized, the entropy value changes discontinuously. Now, for the new framework, this is not necessarily so: whether Boltzmann entropy changes continuously or not, together with the continuous evolution of the micro state, depends on whether the volume of the world actualized changes continuously or not. Showing that it really does lies beyond my abilities. But at least, the new framework doesn't rule out this possibility, by contrast to the old framework, where Boltzmann entropy change must be modelled as discontinuous.

III.11 Conclusion

I have presented the discontinuity problem of the standard framework of BSM, and argued that an alternative way of partitioning phase space into overlapping worlds is called for, given one wants to avoid it: on the hitherto usual method of partitioning phase space into finitely many macro regions of non-zero volume, macro values are always modelled to change in a discontinuous fashion, if they change at all. This is at odds with both the continuous evolution of the micro state, and the expectation that a framework employed for modelling classical thermodynamic phenomena should be able to model continuous change of macro values.

I tried to fix this issue by excising *no-overlap* and *exhaustive-cover* from the old framework. This didn't bring about the results we had hoped for. Removing the latter resulted in "non-regions", which subsequently turned out to be ordinary macro regions after all. But even if one insists that they are not ordinary, the problem of discontinuous change upon the micro state crossing boundaries between macro regions remains. Removing no-overlap, on the other hand, resulted in the micro state instantiating inconsistent macro states at the same time.

Hence, I proposed an alternative framework, featuring overlapping worlds as a replacement for macro regions, and allowing for macro states to be indistinguishable despite being different. The key features of the new framework made it possible to have micro states, residing in the overlaps of worlds, instantiating macro states indistinguishable from macro states that themselves are mutually distinguishable. In this fashion, the macro state of a system can change continuously, without jumps between macro values, and in accordance with the continuous evolution of the micro state. The framework presented, it seems, is indeed a viable alternative to the old one, solving a central, albeit hitherto largely overlooked problem of the latter.

It goes without saying that much more work needs to be done if one really wants to establish the new framework as a superior alternative to the standard framework. So far, it is merely a proposal, sketched in broad strokes. To develop into a full theory, its details need to be worked out. That, however, lies beyond the scope of this thesis. The primary objective I pursued with this chapter was to introduce and develop the concept of overlapping worlds. This will be needed in the following chapter, where I provide an account of how so-called world histories are formed from the discontinuous actualizations of universal particle configurations in Minimal Bohmian Mechanics. There, what I have called the discontinuity problem arises even more severely: how can *discontinuous* actualizations of the configuration of the universe be reconciled with the seemingly *continuous* histories we experience? To do this, I propose to employ the notion of overlapping worlds. My reason to first introduce and develop this concept for BSM is simple: if I want to use a certain concept in order to reconcile an actually discontinuous micro dynamics – and thus an actually discontinuous macro dynamics – with seemingly continuous phenomena, this concept should also be employable to reconcile an actually continuous micro dynamics with what not only seems to be continuous phenomena, but, in classical thermodynamics, are assumed to be continuous phenomena. I have used BSM as a test bed for the concept of overlapping worlds,

to be applied to the case of MBM's discontinuous actualizations in the following chapter.

IV. Worlds in a stochastic universe

In the previous chapter, I have argued that, from the way the usual framework of BSM is set up, the discontinuity problem arises, and that it can be tackled by adopting an alternative framework, allowing for phase space to be carved up into a continuous set of overlapping worlds rather than a finite, discontinuous set of disjoint macro regions. The framework thus construed allows for modelling continuous change of the macro state of a system, in accordance with the continuous evolution of its micro state, and the expectation (from thermodynamics) that the evolution of macro values at least possibly should be continuous.

A similar, despite different problem arises for MBM. Here it is not the continuous evolution of the micro state that is at odds with the unexpected, discontinuous evolution of the macro state. Rather, the evolution of the micro state in MBM, i.e. the evolution of the universal particle configuration Q , is always discontinuous: the configuration randomly jumps through configuration space, thereby instantiating different worlds, one at a time. Successive actualizations of Q at different times are not linked up by (continuous) trajectories. Thus, successive worlds aren't linked up either. On this, Bell's charge against his Everett (?) theory, and likewise against MBM, is based: if worlds – i.e. configurations in his terminology – at different times aren't linked up, and only actualized for an instant, then reference to the past as well as to the future, from within such a world, is empty. Bell (1987b, 135-136): “there is no association of the particular present with any particular past. [...] Everett's replacement of the past by memories is a radical solipsism – extending to the temporal dimension the replacement of everything outside my head by my impressions”.

Bell has a point. It is difficult to explain the common appearance of history evolving consistently over time based on individual, mutually independent, internally static configurations. Barbour, as mentioned earlier, has attempted to provide such an explanation, employing his concept of time capsules. However, as I will argue, it is not easy to reconcile this view with our current best understanding of how appearances arise. Simply put, our current best theories suggest that any experience supervenes on processes happening in the brain, rather than from static, unchanging states. This is at odds with Barbour's proposal, for his time capsules are such static states, and he claims that the experience of history unfolding is due to the brain comparing snapshots from within such static states. To reconcile FAPP-continuous macroscopic change – the history of the universe unfolding consistently – with randomly actualized configurations at the fundamental level of description, is the problem I am going to tackle in this chapter. A suitable link between configurations must be provided, despite their stochastic dynamics. For this, I will employ the notion of worlds in the sense developed for solving the discontinuity problem of BSM.

In the chapter introducing MBM, I have presented the solution Dürr and I have chosen in our recent paper manuscript, based on Barbourian time capsules and FAPP-classical world branches. Neither in this thesis so far, nor in our manuscript has this been spelled out in great detail. In the following, I am going to present my critique of the time-capsule approach, arguing that it is not necessary to account for macroscopic appearances, in particular macroscopic change, from within individual, static configurations. At the same time, I will give a detailed account of how FAPP-continuous, historically consistent world histories can emerge from MBM's stochastic and discontinuous micro dynamics, in accordance with the branching structure of the WF.

IV.1 Worlds from the wave function vs. worlds from the particle configuration

Following the usual presentation in the Many-Worlds community⁶², worlds can be identified with decoherence-induced, quasi-classical branches of the WF. Dürr and I have adopted this notion for our approach. However, to some this approach might look a bit unnatural, given that MBM, as a Bohmian theory, is (at least something similar to) a primitive ontology theory⁶³, with the Bohmian particles as the fundamental building blocks of material objects. Hence, it could be argued, the WF should not play the central role in defining the notion of a world within the MBM framework. Instead, the notion should be defined in terms of the primitive of the theory – the “empirical substructure” in van Fraassen’s (1980, *passim*) terminology.

However, it is to be feared that an argument brought forward against standard BM⁶⁴ carries over to MBM: it can be argued that the particle configuration in standard BM is superfluous theoretical structure. If this holds for MBM, particles and their configurations should not be the basis upon which the notion of worlds is defined, for superfluous structure can be easily excised, thereby rendering said definition of worlds meaningless, and a fortiori undermining the concept that is supposed to account for the appearance of a consistent history in MBM. Let’s assess this argument, headed ‘Many worlds in denial?’, by retracing Valentini’s (2010) and Brown’s (2010) exchange.

IV.1.1 Many worlds in denial?

This argument from the particle configuration being superfluous theoretical structure has to do with the ontological roles ascribed to the two parts of the complete state description in BM. The latter is often read realistically not only with respect to the configuration, but also with respect to the WF. Recall the discussion

⁶² Of which Wallace, in particular with his (2012), is one of the most prominent representatives.

⁶³ With primitive ontology theories being theories that provide a realist reading of some (value definite, usually with respect to position) variable – the so-called primitive variable – of their formalism that refers to the fundamental stuff in physical space or space-time that makes up the world and everything in it. See Allori (2015) for a review of the approach, and my discussion in an earlier chapter.

⁶⁴ Cf. Valentini (2010) and Brown (2010).

of the two-space reading above: because the WF, together with the particle configuration, is part of the complete state description and indispensable for the dynamics as well as the initial distribution of the latter, it should represent something physical. On this basis, the WF is interpreted as “an ontological – that is, physically real – complex-valued field $\Psi(q, t)$ on configuration space” (Valentini 2010, 479).⁶⁵ And because of this being “ontological” of the WF, one could be lead to believe that BM is Many Worlds in denial (ibid.): “If one takes pilot-wave theory seriously as a possible theory of the world, and if one thinks about it properly and carefully, one ought to see that it really contains many worlds – with a superfluous configuration q appended to one of those worlds.”⁶⁶ Even unoccupied branches/wavepackets, according to this argument, should refer to something physical. Take, for example, two wave packets Ψ_1 and Ψ_2 , resulting from a branching process of an original Ψ . They are said to not differ in terms of physicality (ibid., 493): “the motion of Ψ_1 represents a world every bit as bona fide as the world represented by Ψ_2 ”. Valentini (ibid., 495, 497) argues to the contrary: since there really is only one “true history”, i.e. only one wave packet endowed with the particle configuration, and it can be decided within pilot-wave theory which one that is, it is clear that there is only one real world and “the alternative trajectories are mathematical, not ontological.” Hence, following Valentini, BM is not many worlds in denial.

As Brown (2010, 510 ff.) notes in his reply to Valentini, it is not necessary to subscribe to such a “dualist” particle-wf-ontology in the first place. In fact, one might be better off by adopting the Everett interpretation and not buying into the Bohmian particle ontology at all. On that, Brown is right: at least prima facie, one doesn’t have to adopt the Bohmian ontology. However, for our purposes, i.e. for

⁶⁵ Note that Valentini bases his discussion on de Broglie’s original pilot-wave theory, which doesn’t speak of a quantum potential, and in this respect is closer in spirit to what is nowadays called Bohmian Mechanics than Bohm’s original theory (see the distinction above). The latter was formulated by Bohm much later. Most importantly, however, Valentini is committed to an ontologically thick reading of the WF, as opposed to other proponents of BM.

⁶⁶ Note that this and the following quotation isn’t Valentini arguing against BM, but his presentation of the argument brought forward against BM.

answering the question whether the configuration is superfluous in MBM, it is inadmissible to take over Brown's verdict regarding standard BM and straightforwardly apply it to the case of MBM. Also, even if Brown is right and BM indeed is many worlds in denial, such that the particle configuration is rendered superfluous, it is questionable whether this carries over to MBM. Both, BM and MBM, can feature a realist interpretation of the WF. And insofar a realist interpretation of the WF invariably renders the configuration superfluous, MBM falls prey to this objection as well. However, there are different ways in which things can be said to be real. In other words: there are variants of realist interpretations of the WF, and presumably, it depends on the variant whether the argument applies.

So, what precisely does a "realist interpretation" of the WF amount to in BM and MBM, respectively? Valentini likens the WF in BM to the electromagnetic field in Maxwell's formulation of electromagnetism (Valentini 2010, 480): "This provides a particularly good parallel with pilot-wave theory. A charged test particle, placed in an external electromagnetic field $E(x, t)$, $B(x, t)$, will follow a trajectory $x(t)$. One would normally say that the field is real, and that the realized particle trajectory is real; while the alternative particle trajectories (associated with alternative initial positions $x(0)$) are not real." The viability of such a realist reading of the electromagnetic field in Maxwell's formulation of electromagnetism notwithstanding, this suggests that the WF is viewed as being ontologically on a par with the particle configuration: both are regarded as equally real, just like particles and the electromagnetic field are regarded as equally real.

Brown also regards the WF as real, but in a particularly strong sense. He seems to view the WF as a kind of substance, or at least as something sufficiently substance-like, in the sense of a bearer of real, physical properties. Brown & Wallace (2005, 525) emphasize this point in response to Holland: "Quantum mechanics, in our view, has both substance and form before the introduction of hidden corpuscles."

As Brown (2010, 513) writes, in order to avoid the many-worlds-in-denial attack, “what is really needed is an argument to the effect that the wavefunction does not have such potency as the Everettians attribute to it”, with “potency” meaning the WF being able “on its own to account for physical systems, apparatuses, people, etc.”, in other words, to play the role of a substance bearing contingent properties: if the WF alone is substance enough to do all that, the introduction of a further substance in the form of particles indeed seems unnecessary.

It can certainly be denied that the WF *must* be regarded as a substance, or a substance-like entity. Yet, if it is regarded as such, strong arguments are required for maintaining that the particle configuration adds necessary (read: non-superfluous) structure to the theory. Irrespective of that, MBM avoids this objection from the outset. With its particular interpretation of the WF, it doesn’t share the Bohmian WF’s risk of being interpreted as a substance or substance-like: ontologically, the WF and the particles in MBM are *not* on a par. Only the latter play the role of a substance, while the former is regarded as encoding a *dispositional property* of the universe to actualize certain configurations. (Alternatively, adopting a nomological interpretation as discussed above, avoids the problem even more clearly.) Hence, concerns about substance dualism that apply to BM, as discussed by Valentini and Brown, do not carry over to MBM. More to the point, not being a substance itself, the WF alone cannot contain many worlds, or anything, for that matter: unless assuming a realist stance regarding properties without bearers, the WF in MBM is not enough to account for appearances including physical systems, people, etc., given material substance of some sort is required for that. Hence, MBM is not many worlds in denial.

In short: if material substance is required, then either the WF itself is substantial, in which case the configuration is superfluous, or the WF is not substantial, in which case the configuration is not superfluous. Since for MBM, we do require material substance, and do interpret the WF as not being substantial, the configuration is not superfluous. So the objection does not carry over to MBM. The upshot of all

these considerations is that MBM evades the risk of having its basis for the definition of worlds undermined. The fear that this might be the case is unfounded.

IV.1.2 Psychophysical Parallelism

As has just been argued, the substance or substance-like entity in MBM's ontology accounting for appearances is not to be found in the WF, but in the particle configuration. Whether one adopts this view, ultimately, is a matter of interpretive choice: one can instead choose to think of the WF as a substance, thereby rendering the configuration superfluous. Then, one should consider adopting the Everettian view. Both options are viable. But in any case, this decision has some import on how to define worlds within the resulting framework: Everettians must construe worlds from the WF. There is nothing else. In MBM, we are best advised to construe worlds from configurations, for they are the substance. And further down the line, deciding to take either the WF or the configuration as the basis on which to define worlds amounts to deciding for a stance on psychophysical parallelism (PPP): one must decide which part of the ontology of a theory shall be regarded as corresponding to an entity that can figure as an observable and that represents the physical that parallels the psychological in psychophysical parallelism. As Barbour (1999, 226) puts it: "Any scientific theory must establish a postulate of *psychophysical parallelism*: it is necessary to say what elements of the physical theory correspond to actual conscious experience".

In the Everett interpretation, as Barbour points out, the individual branches of the WF fulfil that role. In MBM, on the other hand, worlds defined over configurations fulfil it. Barbour (1999, 302) and Bell agree with this choice: "Now, what did Bell regard as the physical counterpart of psychological experience? Is it in the wave function, as Everett and many others have assumed, or in matter configurations? Bell, like myself, opts for the latter: 'It is ... from the x s [(the configurations)], rather than from ψ , that in this theory we suppose 'observables' to be constructed. It is in

terms of the *x*s that we would define a ‘psycho-physical parallelism’ – if we were pressed to go so far.’ [(Bell 1987a, 128)]”

Let’s dwell a bit on that notion of psychophysical parallelism (PPP). As Heidelberger (2003) shows, PPP can be understood in several ways, ranging from a mere correlation between the physical and the mental that bears no metaphysical weight and allows for a strict substance dualism, to the idea that everything physical has some mental aspect to it (panpsychism). In between lies the idea that the mental is an *aspect of some* physical things, like, e.g., brains (Heidelberger 2003, 238): “The mental and the physical are [...] two different aspects of one and the same entity” (double aspect PPP).

Barbour doesn’t really say where exactly in this spectrum he locates his own account of psychophysical parallelism. On the one hand, he speaks of *correspondence* between elements of the physical theory and conscious experience, which suggests a rather weak form of PPP. On the other hand, it seems clear that he has a specific physical structure in mind, i.e. the brain, the different states of which correspond to experience, thereby suggesting a somewhat stronger PPP (Barbour 1999, 26.): “However, there is increasing evidence that certain mental states and activities are correlated with certain physical states in different specific regions of the brain. This makes it natural to assume, as was done long ago, that there is psychophysical parallelism: conscious states somehow reflect physical states in the brain.” Also (ibid., 266): “It is well known that much processing goes on in the brain and, employing normal temporal language, we can confidently assert that what we seem to experience in one instant is the product of the processing of data [...].”

I will come back later to that issue when discussing its bearings on the construction not only of worlds, but world histories; if one adopts a form of PPP that makes claims stronger than mere correlation between the physical and the mental, and if one considers the findings of modern neuroscience, an account that conceives of

world histories in terms of processes of change between configurations rather than individual configurations alone suggests itself.

But before further discussing PPP and its import on the construal of world histories, we must tackle the question what a world in MBM is supposed to be. After all, world histories are supposed to consist of instantiations of worlds, standing in an appropriate relation to each other. What that relation is will be spelled out later. First, we must account for the relata featuring in said relation. As has been argued, worlds, in the framework of MBM, should best be defined in terms of particle configurations. However, they must not be identified with individual configurations simpliciter. The following paragraph will develop this approach.

IV.2 What is a world?

As mentioned in the introduction to Minimal Bohmian Mechanics, MBM stems from Bell's Everett (?) theory. However, comparing them closely reveals that the two are far from being the same theory, for they differ in an important aspect – their construal of worlds. The just mentioned identification of worlds with individual configurations simpliciter is what Bell had in mind when developing – and immediately rejecting – his Everett (?) theory (Bell 1987a, 133): “Thus instantaneous classical configurations x are supposed to exist, and to be distributed in the comparison class of possible worlds with probability $|\psi|^2$ ”. Note that Bell uses the terms “configuration” and “world” interchangeably: possible worlds – not to be confused with possible worlds as in Lewis' (1986) modal realism – are simply configurations currently not actualized. Actualized configurations, accordingly, are the same as instantiated worlds. In short: for Bell, individual, instantaneous configurations *are* worlds.⁶⁷

⁶⁷ See also his usage of the term at other places in his 1987a: “Let this ‘world’ be simply a large number N of particles” (127); “the initial configuration x is chosen at random from an ensemble of configurations” (129); “a comparison class of possible initial worlds” (129); “a hypothetical ensemble (of worlds!)” (129).

We would be ill-advised to simply adopt Bell's notion of a world, for it is this identification of worlds and individual configurations that leads him to his objection of temporal solipsism, and, a fortiori, to his rejection of the Everett (?) theory. We must be careful that MBM does not fall prey to the same objection. The problem of temporal solipsism has been introduced earlier. It is the reason for Dürr's and my construal of worlds, or rather: world histories, in terms of world branches. To summarize the issue: once worlds are identified with individual configurations as in Bell's account, as soon as the configuration changes, a different world is instantiated. And with this change between individual configurations being discontinuous, worlds are isolated from each other, not only in configuration space, but also in time. On this construal, no connection between worlds can be had. Reference to events future and past – from within a given world – is empty. Worlds, then, have no history in any meaningful sense of the word.

The solution to this problem presented in Dürr & Ehmman (ms) probably isn't quite satisfactory, in that it lacks a detailed description of how our experience of temporal history emerges from instantaneous configurations: our world branch is actualized infinitely many times in each finite time interval, but how exactly the individual actualisations are connected, other than by being in the support of the same world branch, hasn't been made clear yet. Besides, relying too heavily on the branching structure of the WF for the definition of worlds defies the purpose of adopting a Bohmian particle ontology, and doesn't chime with an interpretation of the WF as encoding a dispositional property.

Time capsules, on the other hand, do provide an explanation of our temporal experience. But it can be contested that this explanation is ideal, especially if we aren't forced to assume that there is no external, universal time. Barbour's time-capsule approach will be discussed in greater detail shortly. For now, it shall suffice to say that we want to construe worlds differently.

In order to avoid Bell's criticism, our notion of worlds has to be able to do more than worlds in Bell's sense, identified with individual configurations, can accomplish. At the same time, we must not resort to time capsules or Everettian world branches alone. It goes without saying that, in the end, we want to be able to recover the (quasi-classical) manifest image from the scientific image. As Bell (1987, 125f.) points out, "classical concepts have not [...] been expelled from physics. On the contrary, they remain essential on the 'macroscopic' scale, for '... it is decisive to recognize that, however far the phenomena transcend the scope of classical physical explanation, the account of all evidence must be expressed in classical terms' [(Bohr 1949)]. Thus contemporary theory [employs both quantum wave functions ψ and classical variables x , and a description of any sufficiently large part of the world⁶⁸ involves both". Accordingly, I don't deem it enough for MBM to give an account of the microscopic goings-on. An explanation must be given as to how, at least in principle, macroscopic appearances, including the phenomenon of continuous macroscopic change over time, can emerge from the latter. In other words, we must "save the phenomena" by providing an account of psychophysical parallelism for MBM. Since we do not rely on Barbour's time capsules, individual configurations alone aren't a suitable choice for this purpose.

Instead, we have to construe worlds in a way that allows macroscopic objects to (quasi-)persist over change of configurations, but which also leaves room for their change based on the change of configurations. What we need is a notion of worlds that is able to macroscopically group together configurations that are different microscopically, such that worlds of finite, non-zero volume in configuration space can be construed. In the previous chapter, we have developed such a notion of worlds for the framework of BSM. In what follows, we shall apply it to the case of MBM.

⁶⁸ Presumably, this instance of the term "world" does not refer to a single configuration, but the universe in total.

Just for the sake of argument, let's *preliminarily* define a world in MBM as a collection of configurations generated from "coarse-graining" over possible configurations such that the resulting collection

($\text{WORLD}_{\text{MBM}}$) consists of configurations that are macroscopically *identical*, i.e., that instantiate the *same* macro state of the universe.

Already this preliminary definition avoids the aforementioned problem of Bell's identification of worlds and individual configurations. However, we won't stick with it for long, since it allows for partitioning of configuration space into *disjoint* sets, similar to the partitioning of phase space into *disjoint* macro regions in the usual framework of BSM. Recall that a macro region Γ_{M_i} in standard BSM is a set of micro states $x \in \Gamma$ all instantiating the same macro state M_i , and only M_i . That is to say: Macro regions in standard BSM do not overlap. For reasons extensively discussed earlier, this notion of a macro region as an equivalence class in phase space has been replaced by a definition of "world" for BSM: a world \mathcal{W}_i is the set of all x_j that satisfy SS (i.e., are sufficiently similar in the sense of instantiating indistinguishable macro states) with respect to a given x_i :

$$\forall x_i \exists \mathcal{W}_i \left(\forall x_j (x_j \in \mathcal{W}_i \leftrightarrow x_j \text{SS} x_i) \right)$$

(*world membership*_{BSM})

According to the new definition *world membership*_{BSM}, *all* micro states in a world are required to be sufficiently similar *not to all others*, but *to the central micro state* x_i around which the world \mathcal{W}_i is construed. From this weaker requirement, it follows that not necessarily all $x_j \in \mathcal{W}_i$ are sufficiently similar to each other: SS is a nontransitive relation, and only requiring that all $x_j \in \mathcal{W}_i$ fulfil SS with respect to the central x_i , leaves the possibility open that two other micro states $x_j, x_k \in \mathcal{W}_i$ do not stand in relation SS to each other. At the same time, a micro state x_j being sufficiently similar to a given central micro state x_i and hence belonging to \mathcal{W}_i

around x_i may also be sufficiently similar to another micro state, x_h , say, and hence belong to the world \mathcal{W}_h around x_h .

This new definition of worlds in the case of BSM allowed for overlapping worlds, and hence solved the issue of jumping macro variables despite a continuous micro dynamics. In MBM, we don't have such a continuous micro dynamics. But we still need to recover (FAPP-) continuous macroscopic change. I.e., we need to provide an account of (FAPP-) continuous transitions between different macro states, similar as in BSM. Accordingly, adopting the approach we had developed for BSM and applying it to the case of MBM promises to be fruitful.

Thus, we reuse the definition of worlds we developed for BSM and adapt it for the MBM framework. Some slight modifications are necessary, though. In MBM, we are not dealing with micro states x in $6N$ -dimensional phase space Γ (or, for systems of constant energy, in the $(6N - 1)$ -dimensional energy hypersurface Γ_E), but with configurations Q in $3N$ -dimensional configuration space \mathbb{R}^{3N} .

Accordingly, a world \mathcal{W}_i in MBM shall be defined as the set of all Q_j that satisfy SS (i.e., are sufficiently similar in the sense of instantiating indistinguishable⁶⁹ macro states) with respect to a given Q_i :⁷⁰

⁶⁹ Being "indistinguishable", here, is not necessarily to be understood as being indistinguishable for some external observer or measuring device. The configuration Q usually refers to the configuration of the entire universe, understood as everything there is. By definition, there is no observer or device external to it. In particular, speaking of (in)distinguishability in this context does not imply the existence of a mind for which states are distinguishable or not. Rather, as shall be spelled out in some detail later, saying that macro states are indistinguishable simply means that these states do not amount to a macroscopic difference for a given macroscopic object (MO) within a configuration. Likewise, saying that macro states are distinguishable means that these states do amount to macroscopic difference for the MO. Such a MO can be a mind-bearing observer, but it doesn't have to be.

⁷⁰ Despite being suitable for our present purposes, this definition is still somewhat preliminary, for the following reason: since our definition of worlds via SS hinges on the notion of macroscopic perspective, the meaning of the latter must be made explicit. As we will see, the distinction between the micro- and the macroscopic itself is relative – it differs, depending on the system or object in question. This, we need to account for in our definition. Likewise, we must be more specific regarding the overlap of worlds. Other than for the case of BSM, in order to construe world histories in MBM, a special condition on how worlds overlap – overlap proper – must be met.

$$\forall Q_i \exists \mathcal{W}_i (\forall Q_j (Q_j \in \mathcal{W}_i \leftrightarrow Q_j SS Q_i))$$

(*world membership*_{MBM})

Irrespective of the difference regarding the fundamental space within which the micro dynamics unfolds, and irrespective of the difference in the micro dynamics itself, the feature that worlds defined via SS may overlap, and everything this brings with it, carries over from the case of BSM to the case of MBM.⁷¹

Despite worlds possibly overlapping, *world membership*_{MBM} guarantees – via sufficient similarity, and in accordance with our usual experiences – that macro states instantiated by configurations belonging to a world don't contradict the macro state instantiated by its central configuration, e.g. by there being dinosaurs in Hong Kong in one, but not in the other. In this sense it, suppresses inconsistent macro states: the respective configurations, differing to an extent that they contradict each other on the macroscopic level, certainly are macroscopically distinguishable. Hence, SS does not put them in each other's worlds.

Consequentially, the application of SS leads to individuation (separation, “branching”) of worlds whenever the grouping of configurations into a world would lead to inconsistencies with its central configuration.⁷² But note that the requirement of sufficient similarity does not suppress the grouping of microscopically inconsistent configurations into one world: such configurations, differing only microscopically while being macroscopically indistinguishable from a central configuration, still belong to the world of the latter. This is to say that *world membership*_{MBM} also guarantees that individuation of worlds in this sense only occurs when necessary. Preventing such “over-individuation”, SS avoids the problematic identification of configurations and worlds: with a microscopic difference being sufficient for the individuation of worlds, every world would contain only one configuration, since every configuration differs from every other

⁷¹ Recall from above that SS is reflexive, symmetric and nontransitive.

⁷² Cf. Barbour's (1994, 408) similar requirement for records contained in time capsules: “A time capsule [...] appears to contain records of the past, and these records are mutually consistent”.

configuration microscopically. Accordingly, worlds could not overlap, which ultimately leads to temporal solipsism. For example, on such an account, it is utterly unclear how the setup of an experiment and the recording of its results hang together (unless we assume that all of this emerges as an illusion from within one Barbourian time capsule). It is by avoiding over-individuation that MBM evades temporal solipsism.

IV.3 Superpositions in MBM

One should be careful, though, to not confuse the grouping of microscopically inconsistent configurations with a superposition. Superpositions are a feature of the universal WF and its dynamics. It evolves independently of the actualized particle configurations, and likewise of their grouping into worlds. If it contains superpositions, it does so irrespective of how configuration space is coarse-grained into worlds, and which macro states are instantiated.

Having said this, it might be helpful to see how MBM accounts for interference experiments probing superpositions by looking at an example. This might also be useful to render the similarities and differences between MBM and some of its main competitors a bit clearer. Recall from the introduction to MBM my presentation of the measurement problem, following Maudlin (1995). One can categorise different quantum theories and/or interpretations thereof based on which route they take for the solution of this problem in dealing with the general presence of superpositions. In the following, before coming to MBM, I present one example from each of these categories: Copenhagen, not accepting the unaltered Schrödinger dynamics and modifying it by adding a collapse postulate; Everett, replacing (universal) value indefiniteness by value multiplicity (cf. Wallace 2012) for “macroscopic superpositions”; and BM, denying completeness of the WF, adding Bohmian particles/configurations to complete the state description. From this categorisation, we recover MBM’s heritage: it is clearly to be grouped with BM, in that it adds particles/configurations to complete the state, rendering the complete

state value definite at any given time while accepting the unaltered Schrödinger dynamics. Yet, despite being value definite at any given time, MBM's configuration is not determined by configurations at other times, for it lacks the deterministic particle dynamics, i.e. the GE, of standard BM. Maudlin (*ibid.*, 13) refers to this as the "problem of effect". MBM solves it by appeal to worlds and world histories. Having to make this appeal is the price to pay for the excision of the deterministic GE. It results in the fact that, while still being value definite at any given time, different world histories, carving out the branching structure of the universal WF, are instantiated at random, rendering the theory Everettian to the extent that none of the branches of the universal WF shall be forgotten – all of them are actualized again and again, although not simultaneously. In this restricted sense, MBM inherits value multiplicity from Everett.⁷³

Take the usual single particle/wave packet double-slit experiment: A particle/wave packet source that can be set up such that individual particles/wave packets are emitted, one at a time; a double slit the particle/wave packet is passing through; and a screen for measuring the particle/wave packet position behind the double slit. Despite the system generally being in a superposition, we do not observe a macroscopic superposition at the screen.⁷⁴ Now, different theories/interpretations tell different stories of what is going on when this experiment is carried out.

According to Copenhagen, the wave packet, travelling through the apparatus, is in a superposition, until, upon measurement, it collapses onto definite position state, with probabilities according to BR. What counts as a measurement, and accordingly, where and when collapse occurs, depends on how we define and divide the

⁷³ In the less restricted sense as used by the Everett community, MBM does not feature value multiplicity: at each instant, only one configuration is actual, thus only one branch or world history actualized. Thus, at each instant, there is no multiplicity of values in the Everettian sense. Rather, multiplicity in the restricted sense appropriate for MBM appears only when temporally coarse graining over finite intervals of time.

⁷⁴ In a single event, that is. Of course, we are "observing" macroscopic superposition, in the sense that in a sufficiently large ensemble of single events, events approximately are distributed with Born statistics over the screen. That is, after many individual runs of the experiment, an interference pattern emerges on the screen.

quantum system to be investigated from its environment, construed as classical. An experimenter in the lab, e.g., would place the measuring device, in this case the screen, on the opposite side of this so-called Heisenberg cut than the system to be measured. A colleague of the experimenter, however, might place the cut after the first experimenter, thereby treating the latter as part of the system under investigation, and so on.⁷⁵ This is of importance for our considerations: placing the Heisenberg cut decides on where and when the collapse of the WF happens. And it is the collapse of the WF that is supposed to account for the quantum-to-classical transition in Copenhagen, in the sense that it accounts for the fact that we do not observe a macroscopic superposition in single events.

By contrast to Copenhagen, Everett does not resort to a collapse postulate in order to explain the absence of macroscopic superpositions from our observations. Rather, decoherence accounts for the vanishing of interference terms, resulting in the branching structure of the WF, i.e. FAPP-independent, quasi-classical world branches all of which are realized simultaneously. As is regularly stated, decoherence alone provides no solution to the measurement problem, for it does not do away with value indefiniteness. Instead, Everett's solution to the measurement problem is the replacement of value indefiniteness by value multiplicity for macroscopic superpositions. Universally, superpositions remain. But as inhabitants of one such emergent world, we do not encounter macroscopic interference, despite superpositions still being present in the universal WF. As Wallace (2012, 36f.) puts it: "this, in a nutshell, is what the Everett interpretation claims about macroscopic quantum superpositions: they are just states of the world in which more than one macroscopically definite thing is happening at once. Macroscopic superpositions do not describe indefiniteness, they describe multiplicity". And later (ibid., 38): "If the correct way to understand such superpositions is as some sort of multiplicity, then our failure to observe that multiplicity is explained quite simply by the fact that we live in one of the 'worlds'

⁷⁵ Cf. e.g. Schlosshauer & Camilleri (2011) and references therein.

and the other ones don't interact with ours strongly enough for us to detect them." Note that branching via decoherence depends on "magnifying" microscopic superpositions (ibid., 99): "branching is caused by any process which magnifies microscopic superpositions up to the level where decoherence kicks in". Everettian worlds, thus, are an emergent feature of a higher, macroscopic level of description, just like worlds in MBM are. Here, too, the distinction between the micro- and the macroscopic is rather vague (ibid, 100): "there is no 'finest' choice of branching structure: as we fine-grain our decoherent history space, we will eventually reach a point where interference between branches ceases to be negligible, but there is no precise point where this occurs".

Standard BM also doesn't resort to WF collapse, although the notion of effective collapse is common. This does not refer to a collapse of the universal WF as in Copenhagen, however. It basically means ignoring those branches of the WF that, after decoherence, are not effective in guiding the particle on its trajectory, for the interference term being vanishingly small. That branch of the WF in the support of which the particle resides is then called the effective WF. But again, the universal WF does not collapse in BM. As described in Dürr & Teufel (2009, 180), "this collapse is not a physical process, but an act of convenience. It is introduced because it would simply be uneconomical to keep the ineffective wave functions, and the price we pay for forgetting them amounts to nothing". Here, too, the superposition in the universal WF remains, but the particle selects one branch that is realized. Which branch that is – unbeknownst to the experimenter, according to the usual reading of the QEH – is already determined at the preparation of the experiment, i.e. in the particle source. Over the course of many repetitions, slight deviations in the preparation correspond to different branches being effective, such that different trajectories are realized, resulting in an interference pattern to appear on the screen.

In MBM, things are a bit different. The QEH, here, does not provide a measure of ignorance about the precise preparation of the particle. Rather, it provides a

stochastic guiding law for the actualization of configurations. As we have seen above, SS groups different configurations, amounting to a microscopic superposition but being sufficiently similar to a central configuration macroscopically, into the same world, as long as the superposition has not been blown up to macroscopicity by resulting in a macroscopic effect. Decoherence in MBM results in branching of the WF, just as it does in Everett and standard BM. But neither are these branches all realized simultaneously, nor are any of them effectively discarded. Let's denote the configuration with the particle residing at a position within the left slit Q_l , and the configuration with the particle residing at a position within the right slit Q_r . These do not yet correspond to different branches being actualized/distinguishable macro states being instantiated, because they are macroscopically indistinguishable: via SS, $Q_r \in \mathcal{W}_l$ and $Q_l \in \mathcal{W}_r$. Compare the presentation of the Everettian account of this situation above and in particular, note the similar appeal to the microscopic-macroscopic-distinction: the quantum state is not yet decohered. No magnification of the microscopic superposition has taken place so far. Hence, although the notions are vague, it is clear that, in the Everettian terminology, branching has not yet occurred. Likewise, in MBM, worlds have not yet been individuated, i.e. configurations still lie in each other's worlds: with the particles residing at different positions, but both configurations being sufficiently similar to each other, their worlds "contain" a microscopic superposition. That is, these worlds overlap⁷⁶: via SS, $Q_r \in \mathcal{W}_l$ and $Q_l \in \mathcal{W}_r$. It is only at a later stage of the experiment, with different macroscopic points on the screen lighting up, that configurations $Q_{l'}$ and $Q_{r'}$ are macroscopically distinguishable⁷⁷. Other than Q_l and Q_r , the former are not sufficiently similar, i.e. do not instantiate indistinguishable macro states. Thus, $Q_{r'} \notin \mathcal{W}_l$ and $Q_{l'} \notin \mathcal{W}_r$. It is at this point where, in the Everettian picture, branching takes place. We see that it is not the presence of a superposition per se, but the presence of a *macroscopic* inconsistency, that individuates worlds, both in Everett and MBM. In particular,

⁷⁶ As will become clear later, they not only overlap, but overlap *properly*.

⁷⁷ A reminder: on MBM, configurations at different times are stochastically independent. Hence, one must not think that $Q_{l'}$ somehow follows from Q_l . That is, it is not possible to infer from the position at which the screen lights up the position of the particle at an earlier time.

microscopic superpositions alone have no bearing on the individuation of worlds. Worlds, and world histories, are thus an emergent feature, like branches in Everett. In fact, world histories “trace” Everettian branches. Just like branching in Everett (cf. Wallace 2012, 3.11), the individuation of worlds depends on what we take to be sufficiently similar on the macroscopic level.

By now, I already have used the terms “microscopic” and “macroscopic” extensively without making particularly clear what they are supposed to mean. Let me make up for that in the following paragraph.

IV.4 The Microscopic-Macroscopic-Distinction and its import on the construal of worlds

It should have become evident by now that the distinction between the micro- and macroscopic is not very precise. Worlds, construed via a relation of sufficient similarity, spelled out in terms of macroscopic indistinguishability, thus are quite pragmatic and relative. They are pragmatic, first and foremost, because it often depends on practical limitations between which states we can distinguish, and between which we can't. In the chapter on BSM, I have elaborated that macroscopic (in)distinguishability is observer- or instrument-relative: whether two macro states can be told apart or not depends on “resolution power” – the abilities and limitations of our sensory apparatus, and/or on the precision of the instruments we use. In this sense, macro states are distinguishable (or not) relative_{observer} to a given observer and/or instrument with a certain resolution power.

(In)distinguishability also depends on our point of view: from far away, two states might be indistinguishable that are distinguishable at close distance. Consider two pixels on a computer screen. If one is lighting up, the other is dark, and vice versa. Say that, very close to the screen, you are able to tell these states apart. But from a certain distance and beyond, you aren't. Your sensory apparatus doesn't change, in

terms of its abilities (in this case, quite literally the resolution power of your eyes), depending on your distance from the screen. But the *apparent* separation between two points is smaller from afar than it is from close by. So at some distance, your resolution power doesn't suffice anymore to resolve the spatial distance between the (middle points) of the two pixels. So indeed, it seems that distinguishability does not only depend on resolution power, but also on perspective.

Think of everyday phenomena: how far must the moon have wandered over the night sky for you to recognize that it now is somewhere different from where it was before? How about a snail? The perceived angular velocity of the moon, roughly being the same as the true angular velocity of the earth's rotation around its own axis, is about $360^\circ/\text{day}$. Helix pomatia, at a distance of one metre, and with a speed of 4.2 metres/hour, can run 16 full circles around you in the time the moon completes one circle. Its perceived angular velocity, from your perspective, is about 16 times as high as the moon's.

With distinguishability not only depending on resolution power, but also on perspective, the construal of worlds, making use of distinguishability, also depends on perspective: configurations that, from your perspective, are distinguishable, and hence don't belong to each other's worlds, might not be distinguishable from the perspective of someone else, or, in general, from the perspective of another being, or system.⁷⁸ This other being might be a million light years away, or as big as the solar system, and hence not being affected by the difference between, for example, your cat sitting on the couch or sleeping on your bed, just like you are not affected by the difference between two slightly different molecular compositions of the

⁷⁸ As pointed out above, the notion of macroscopic (in)distinguishability is not reserved for sentient beings (consciously) observing differences between macro states. In a broader sense, adequate for our treatment of MBM, two states being macroscopically distinguishable simply means that it makes a macroscopic difference for a given macroscopic object (MO) which one is instantiated, and likewise, two states being macroscopically indistinguishable simply means that it makes no macroscopic difference for a given MO which one is instantiated. Such MOs may very well be sentient beings, but they don't have to be. Speaking of "(in)distinguishability" is a convenient *façon de parler* that should not be taken to imply that a (conscious) observer or, in general, a mind of some sort or other is of fundamental importance.

beast. It is here where the relativity of the terms “microscopic” and “macroscopic” enters the scene. It is a matter of perspective, not least of size, what does and does not count as micro- or macroscopic. Hence, the usual construal of the microscopic-macroscopic-distinction as two levels of description, appropriate for two different length scales defined in absolute terms, is too simple.

Let’s elaborate on this. Length is a measure usually defined over the positive reals (see e.g. Tao 2011), such that lengths, in general, can have any positive real-valued magnitude. The definition of a standard unit, which, for length, is the metre, is a pragmatic convention. And so is picking out a certain interval of the reals, measuring it in terms of some standard unit, and call it “the macroscopic length scale” or “the macroscopic realm”. The same goes for the microscopic. Especially in descriptions like “the scale of objects that can no longer be seen by the naked eye” of the microscopic scale, the pragmatic nature of such assignments shows. See, for example, the definition in Reif (1965, 2): “It is useful to introduce a distinction between the sizes of systems whose description may be of interest. We shall call a system “*microscopic*” (i.e. “*small scale*”) if it is roughly of atomic dimensions or smaller (say of the order of 10 \AA or less). For example, the system might be a molecule. On the other hand, we shall call a system “*macroscopic*” (i.e., “*large scale*”) when it is large enough to be visible in the ordinary sense (say greater than 1 micron, so that it can at least be observed with a microscope using ordinary light).”⁷⁹

Definitions like this are not only pragmatic, but highly dependent on the contingencies of one part of the human sensory apparatus. But there are other senses, and modes of sensation, in humans as well as in other beings with different sensory abilities. In general, different systems might only be affected by objects (or events, or processes) at greater scales or already be affected by objects (or events, or processes) at smaller scales than we as humans typically are. This suggests,

⁷⁹ Jaeger (2014, 897) also cites this definition and comments: “The appeal of such a notion is not surprising, given the dominance of the visual sense in humans”.

instead of defining the micro- and macroscopic relative to the human sensory apparatus, that we define these notions relative to any two systems in general. After all, we are interested in describing systems and their relations to each other, whether humans are involved or not. And by doing so, we don't lose anything. What we hitherto have called microscopic, with respect to our sensory abilities, can still be called microscopic, as a result of the relation that holds between the system we call "us" and the system to be described. On the other hand, though, we gain the possibility of using the microscopic-macroscopic-distinction in other cases not involving humans or other beings of a certain size. What will matter are only relative measures and how things (objects/events/processes) at different scales affect each other.

One might be tempted, for MBM, to simply adopt a distinction between the micro- and the macroscopic based on decoherence⁸⁰: a system, then, can be called macroscopic once it contains sufficiently many degrees of freedom, or is large/heavy enough such that "the environment [...] continually interacts with, and becomes entangled with, that system" (Wallace 2012, 77). Under such circumstances, systems decohere in a sufficiently short time interval, such that interference becomes practically irrelevant. However, such a construal does not evade vagueness. Recall (ibid., 100): "there is no precise point where this occurs". More importantly though, decoherence is a process that accounts for the branching structure of the WF. This is fine for Everett, but not enough for our purposes. With the particle configuration as MBM's referent and basis for its psychophysical parallelism, our micro-macro-distinction should be spelled out in terms of objects and their constituents as being part of the universal configuration.

Let us therefore adopt the following characterization of the microscopic-macroscopic-distinction: a scale is called microscopic with respect to another, called macroscopic, if change on the level of the former doesn't necessarily imply change

⁸⁰ Cf. also the discussion in the next paragraph.

on the level of the latter, i.e., if it is possible that, from the perspective of the macroscopic object, nothing has changed despite actual changes happening on the level of the microscopic.⁸¹

Through this characterization, it becomes clear that, although the distinction is dependent on the relation of levels, the relation itself is objective: Whether it makes any difference if your cat sits on the bed or on the sofa, is relative to the level from which the situation is described. From your perspective, it does make a difference. Hence, the difference in the cat's position, described from the level of your perspective, is regarded as macroscopically distinguishable. Changing perspective to a relatively large scale can change that status. From the perspective of the solar system and the level used to describe it, the change in your cat's position is irrelevant, hence regarded as microscopic. Yet, despite this relative dependency on the level of description, it is objectively so that, from your perspective, the position of your cat is relevant, while it is not from the perspective of the solar system.

IV.4.1 The Microscopic and the Macroscopic Domain vs. the Quantum and the Classical Domain

A characterization like the one just suggested can also help to prevent the conflation of the “quantum scale” or “quantum domain” with the microscopic, and of the “classical scale” or “classical domain” with the macroscopic, respectively. Not everything small is quantum, not everything big classical. Rather, we should use the terms quantum and classical for phenomena best described by the respective theories, independently of the size of the objects involved. Jaeger (2014, 896f.) discusses this topic (and traces the history of the use of the term macroscopic):

⁸¹ Note the resemblance of the micro-macro-distinction adopted for MBM, and for BSM (cf. the previous chapter), respectively. This is a recurring theme, arguably rendering MBM – as well as other particle-based quantum theories adopting the notion of worlds as collections of sufficiently similar micro configurations – especially suitable for the replacement of BSM in cases where a quantum rather than classical treatment is appropriate (see the following paragraph).

“Related to the popularity of this approach [(i.e. the Copenhagen Interpretation)] is the widespread assumption, which is often made without circumspection and most often made in popular discussions, that the macroscopic realm and the regime of classical physics are one and the same, sometimes resulting in the use of the terms classical and macroscopic interchangeably. This has occasionally been noted, for example, by Żurek in his well known discussion of the attempt to understand the transition from quantum-mechanical to classical-mechanical behavior through the notion of decoherence[:] ‘In the absence of a crisp criterion to distinguish between quantum and classical, an identification of the classical with the macroscopic has often been tentatively accepted. The inadequacy of this approach has become apparent as a result of recent developments: A cryogenic version of the Weber bar – a gravity-wave detector – must be treated as a quantum harmonic oscillator even though it may weigh a ton...’ [(Zurek 2002)]” Later, Jaeger (2014, 898) goes on: “To reach the identification of the macroscopic with the classical, one would need a system being macroscopic as both necessary and sufficient for its being classical. As noted in the introduction, however, some quite massive gravity-wave detectors are required to be treated quantum mechanically. Moreover, a system’s having large quantum numbers, its having large mass, and its having a large spatial extent do not always coincide: one example of this would be that of a compact astronomical object which is of “microscopic” spatial extent while also being considered “macroscopic” by virtue of a large mass – having a Schwarzschild radius of a size larger than an atomic nucleus but smaller than an atom.”

These latter remarks suggest that what we have said above about length holds analogously for other measures. Defining intervals over the space of possible values and calling them the microscopic or macroscopic scale, first of all, is a matter of pragmatic convention, usually based on the human senses and everyday experiences, which are a contingent matter of fact. That doesn’t change when considering other measures instead of length, e.g. mass. A definition based on relations between systems can reproduce the common notions of micro- and macroscopicity by relating the system of interest to systems roughly of the size,

mass, or what have you we encounter in our everyday experience, which can be, if wished for, a human being. But it can also reproduce many more, for example by relating systems of the type just described to other, much bigger (or heavier) systems, a human being to the solar system, say. Suddenly, objects we encounter on a daily basis, including other human beings and ourselves, are rendered microscopic with respect to the macroscopic solar system.⁸² As a consequence, it becomes almost trivial to explain why it would be wrong to conflate the domains thus variably rendered micro- or macroscopic with the domains of quantum or classical mechanics, despite the fact that *often* (not always!), the domain where quantum mechanics is adequately used for the description of phenomena is one of the small, light, etc., and the domain where classical mechanics is adequately used is one of the big, heavy, and so on, with “small”, “light”, “big” and “heavy” as judged from the human perspective.

So much for now about the microscopic-macroscopic-distinction. Earlier, I promised to come back to the issue of psychophysical parallelism. My claim was that an account of world histories in terms of processes of change between configurations will suggest itself, provided we adopt a more than minimal version of psychophysical parallelism and give some credence to the best neuroscientific theories we currently have. To first address these requirements, we shall make a short detour, discussing *a minor issue with Barbour’s account*.

IV.5 A minor issue with Barbour’s account

The issue I identify in Barbour’s account is not a theoretical defect *sensu stricto*. Rather, what I find somewhat problematic is him ignoring, against what seems to be his own conviction, the important role that processes, and hence change – as opposed to states, individual events or static objects – play when it comes to

⁸² Which seems to be an advantage. It can be regarded a theoretical virtue to exclude reference to contingencies like the usual size of humans from general descriptions, and another virtue to be able to reintroduce them as special cases when of interest.

psychological experience and its physiological basis⁸³, and not providing a satisfying answer to the question of how, in his account, conscious experience arises from static configurations. Processes might not be metaphysically – let alone logically – necessary for experience, but modern neuroscience suggests that, at least for us, physiologically they are. Given we can trust our best theories, we should believe that experience arises from processes, not static states.⁸⁴ There’s even a measurable minimum amount of time required – in our context, this translates to a minimum amount of change – in order for consciousness to emerge from an active group of neurons.⁸⁵ We should not sweep this evidence under the carpet lightly, especially not if we don’t have to.

Over the course of more than two decades, Barbour in his writings⁸⁶ develops a timeless cosmology, including an account of psychophysical parallelism. It bears some resemblance, but is far from identical, to the account presented in this thesis. As has been pointed out before, it is similar in that it features instantaneous configurations at the heart of its ontology, and as the physical basis of psychological phenomena: Barbour (1999, 302), following Bell, regards particle configurations as “the physical counterpart of psychological experience”. However, contrasting my account, he deems individual configurations – albeit somewhat special configurations, i.e. time capsules – sufficient to give rise to psychological experience. In this respect, Barbour’s account resembles Bell’s Everett (?) theory. What’s more, he not only conceives of worlds as individual configurations (like Bell did), but also as eternal, static objects. This, I hold, gives rise to concern when it comes to psychophysical parallelism, understood in a more than minimal sense.

⁸³ Cf. (Barbour 1999, 266): “it is well known that much processing goes on in the brain and, employing normal temporal language, we can confidently assert that what we seem to experience in one instant is the product of the processing of data coming from a finite span of time”.

⁸⁴ Despite the common usage of the term “brain state”, especially in the philosophy of mind. Such “brain states” are not states in the strict sense of an individual, unchanging micro configuration of the brain.

⁸⁵ See, e.g., Sacks (2017, 173ff.)

⁸⁶ Cf. Barbour (1994), Barbour (1999), Barbour (2009) and his recent, unpublished book manuscript *The Janus Point*.

Intuitively, one should think: experience, in particular experience of motion, requires change of configurations.⁸⁷ Change in the configuration of what I experience as moving – the Kingfisher in Barbour’s example⁸⁸ – and, at least as importantly, change in myself that reflects the change of the object I observe and allows me to experience that change. Both, change in the object as well as change in the subject – me – require change in the configuration of the universe. So, experience, all the more experience of motion, should not occur within individual, static configurations. Yet, Barbour in his account construes experience as being had from within such configurations. So, given that experience indeed requires change of configurations, Barbour’s account is misguided.

For our purposes, we can set aside the question of how exactly change in the experiencing subject is supposed to be linked up with change of an external object, such as to adequately reflect the latter’s movement. Already the case of experience simpliciter is enough to convey the sense in which Barbour’s account is problematic. Given there are subjects that have experiences of some sort. Since these subjects are part of the universe they live in, and experience seems to require change in that which has the experience, it seems to be the case that the configuration of the universe should change in order for experiencing subjects to arise. But Barbour’s account doesn’t allow for change of the configuration. This poses a problem Barbour tries to circumvent by denying at least one of the premises that 1) experience requires change within the experiencing subject and 2) change within an experiencing subject requires change in the configuration of the universe. I doubt that he would object to the second premise: since the experiencing subject is part of the configuration of the universe, every change within the former amounts to change within the latter, if only on the microscopic level. Instead, he denies the intuitive premise 1) and introduces the notion of *time capsules*: individual, static

⁸⁷ Note that this really is just an intuition. Intuitions can be wrong, so without further investigation, they never suffice to reject an account. Accordingly, it would be premature to reject Barbour’s account just because it doesn’t chime with our intuitions. But they can very well guide us in asking the questions that need answering. Personally, I deem it useful to let intuitions tell me where to look, but not what to find.

⁸⁸ Barbour (1999), passim.

configurations that contain what seem to be records of the past. “Somehow”, by comparing the parts of the configuration that are regarded as records, the illusion of time flowing is said to arise. Let’s take the existence of such time capsules for granted, if only for the sake of the argument. Does this solve the problem? I hold that it doesn’t, even though at first glance, one might think that time capsules are sufficiently structured to account for experience. But the problem of *experience without change* is merely pushed back a little. It can now be formulated as follows: given there are time capsules, how is the operation of comparing several records contained in a time capsule supposed to work from within such an individual, static configuration? Doesn’t such an operation itself require a process to supervene on? Tailored to the case of human experience: how does a brain compare the records contained within a configuration while itself being caught in a static state, i.e. without itself undergoing change, i.e. without brain processes happening?

Recall that Barbour advocates psychophysical parallelism: experience should somehow be correlated with something physical. And indeed, he does name this something – time capsules. Accordingly, my objection to Barbour’s account is not that he doesn’t meet his own demand to “establish a postulate of psychophysical parallelism”. He clearly does that, and from a purely metaphysical or logical perspective, there is no problem with his proposal. It could well be the case that this “something physical” with which experience corresponds is a single, static state: “Above” each time capsule, a mind might “hover” and hallucinate change in accordance with the records present in its associated time capsule. Consequentially, Barbour’s account does work, under the assumption of such a minimal form of PPP, featuring nothing but this unexplained correlation between mind and configuration. This amounts to a dualist form of PPP: other than the thin correlation between configuration and mind, nothing accounts for them hanging together. Now, my worries are that this minimalist PPP is not supported by our current best understanding of how the brain gives rise to experiences, and that, on top of that, even Barbour himself might envisage a stronger account. But irrespective of that latter point, in light of our best theories, one should at the very least provide some

good arguments in favour of such an assumption, especially when putting forward a theory that so strongly hinges on PPP. Barbour does not provide such arguments. To the contrary, his remarks about the relation between brain processes and psychological experience even suggest that his idea of how psychophysical parallelism works might be far from minimal (ibid., 266): “it is well known that much processing goes on in the brain and, employing normal temporal language, we can confidently assert that what we seem to experience in one instant is the product of the processing of data coming from a finite span of time”.

Charitably assuming that such a *product* of data processing can be fully present and experienced in a single instant, the question remains how the *processing* bringing about this product is supposed to happen within an unchanging configuration. That would amount to a process that, qua definition of what a process is, is not a process. Barbour, by resorting to explanations as quoted, at least implicitly seems to acknowledge that an account featuring more than a minimal form of PPP must allow for change of configurations. His qualification “employing normal temporal language” does not resolve the issue. Understandably, he resorts to such formulations, indicating that his talk of processes is merely a *façon de parler*, not to be taken literally, for there are no literal processes within static configurations. But then, he should give an alternative account of how experiences arise, in particular, how the brain is able to compare and put into order the several snapshots it contains, without real processes happening. But he doesn’t provide such an account. All we get (ibid., 266f.) is “the working conjecture [...] that when we think we see motion at some instant, the underlying reality is that our brain at that instant contains data corresponding to several different positions of the object perceived to be in motion. My brain contains, at any one instant, several ‘snapshots’ at once. The brain, through the way in which it presents data to consciousness, somehow ‘plays the movie’ for me.”

This sounds rather suspicious. I have no objection towards the idea that the brain can at one instant contain several records or ‘snapshots’ of an external object, each

representing the latter in a different state. What I find questionable, in addition to the qualms already discussed, is the following: the brain in this account is presenting data to consciousness/“plays the movie” to the experiencing subject. This, taken literally, again amounts to a mind-body dualism that is at least not supported, despite not logically being ruled out, by Barbour’s more than minimal psychophysical parallelism as reconstructed above. Last but not least, all this is supposed to happen “somehow”, which is a rather hand waving explanation, if it is an explanation at all. All things considered, it remains utterly unclear how experiences are supposed to emerge from a static configuration.

The problem of Barbour’s account seems to be condensable into the following argument: in order to account for the illusion of change, he requires an entity – a brain – capable of comparing snapshots from within a static configuration. At the same time, it stands to reason that this entity can only fulfil this operation if it itself undergoes change, i.e., if brain processes can happen. But in a single, static configuration, there are no brain processes, and thus no entity capable of comparing snapshots from within it. Hence, there is no account for the illusion of change. Barbour’s account seems to block itself.

Granted, this argument works only if one doesn’t assume a mind strictly independent of any physical supervenience basis. Such an entity could carry out the necessary operation of comparing records contained in time capsules. But Barbour obviously doesn’t make this assumption, given his requirement for PPP with configurations as the basis. After all, assuming an independent mind would render the whole effort moot: richly structured time capsules as the physical aspect of PPP, the comparison of records from which the illusion of time emerges, all this is superfluous if we simply assume some entity capable of having experiences independently of the physical basis. Since Barbour deems it necessary to go through all this trouble, I deem it safe to say that an entity experiencing independently of a physical basis is not what he envisages for his account. Following his own presentation, it is the time capsule itself, via its rich internal structure, and not

some independent entity, that accounts for experience (ibid., 267, my emphasis): “A time capsule [...] is so highly structured that *it creates the impression of motion.*” Accordingly, I still see no way out of the predicament Barbour’s account is in. On the one hand, he doesn’t allow for change of configurations, because all configurations are supposed to be actualized eternally. On the other, he arguably assumes a more than minimal form of PPP that is too demanding to work with an account that doesn’t allow for change of configurations. Of course, we must consider the possibilities that a) I have misinterpreted Barbour with respect to his stance on PPP and he in fact deems a minimal PPP with its dualist consequences acceptable; and that b) I have read him correctly and he had a more than minimal PPP in mind, but would, in light of the considerations above, reconsider and adopt a minimal PPP. In both cases, I have no objection against his account. Alas, I would not adopt it for MBM, if only for the simple reason that for MBM, we don’t have to buy into a minimal PPP and, a fortiori, dualism: MBM allows for change of configurations, so we may as well try to come up with an account featuring a more than minimal form of PPP, one that doesn’t throw overboard our current understanding of how the physical and mental hang together.

In light of all that has been said so far, I suggest considering a theory in which psychophysical parallelism is stronger than a mere claim of correlation while being compatible with the underlying ontology and dynamics (and with our current best theories in neurology and cognitive science). To be precise, we should consider an account that allows for change of configurations, in the sense of different configurations being actualized at different times, such that experience can arise from processes of change rather than from a static entity like time capsules. This is not a rejection of Barbour’s account *as a whole*: the universal WF (or its replacement) still could be time-independent. The relevant difference between Barbour’s account of the emergence of conscious experience and mine is that I don’t assume conscious experience without physical change. So let’s see whether we can come up with an account in which experience stems from processes, and that is compatible with MBM’s discontinuous actualizations of configurations. The

gist of such an account has already been provided, in the form of our definition of worlds – via *world membership*_{MBM} – above. Now, as announced earlier, we need to expand on this definition, and then distinguish between overlap simpliciter and overlap proper. This will allow us to construe (FAPP-) continuous world histories from successions of individual worlds, instantiated by individual configurations.

IV.6 The macroscopic object

The first problem that poses itself when developing a theory that construes experience as emerging from processes rather than static configurations can be stated like this: the universe has to know how to change configurations (and choose their order) by itself. We cannot rely on a mind to do so. Here is the reason: as suggested by our current neurological understanding, cognition, experience, etc., in general what we would call “mind”, correspond to brain processes, rather than brain states *sensu stricto*, understood as single, instantaneous configurations. These brain processes, as a series of configurations/changes of configurations, occur independently of what emerges from them, in the sense that the “emergent” feature is not causally affecting its “basis”.⁸⁹ What has to happen in the universe, all by itself and without a mind making it so, is a series of actualized configurations that forms a consistent history without macroscopic “jumps”, but that allows for macroscopic change over time. Consecutive actualizations of configurations within such a history must resemble each other so closely that, macroscopically, i.e. for a given macroscopic object within those configurations, the actualization of one after the other does not amount to macroscopic change, but actualizations further apart temporally may amount to macroscopic change. Only then (which is not to be understood as temporal order, but as a conditional conjunctive) can minds emerge,

⁸⁹ The reader may think of this as a form of the principle of causal closure of the physical, as it is often named (cf. Kim 1993), although in MBM, individual actualizations of configurations, being stochastically independent, are not causally connected. The principle in the context of MBM hence does not entail that there are causal connections between physical events on the fundamental level. It simply means that there are no causal connections between the mental and the physical, in particular, that there isn’t anything non-physical causing actualizations and/or grouping them together into worlds, in accordance with the double aspect form of PPP we adopt.

given the emergence of minds requires processes to happen. Following Heidelberger (2003, 236), psychophysical parallelism in the sense of a “double aspect theory”, where “the mental and the physical are [understood as] two different aspects of one and the same entity” is incompatible with the idea that the mind causes the process to happen. Conversely, it is incompatible with the idea that the process causes the mind to be (ibid., 238f.): “[Double aspect PPP] is defined as noncausal and therefore noninteractionist. But this noncausal interpretation [...] results from the definition of the psychical and the physical in terms of the perspective in which something is given. Viewing the physical as something that causes the mental, or vice versa, results from scrambling differing perspectives. [...] We can demonstrate that distinguishing perspectives is nothing mysterious by considering a bent coin. It would be ridiculous to say that a dent on the head's side causes a bulge on the tail's side. While both sides of the coin are intimately connected, their joint occurrence has nothing to do with causality”.

We now have two options. Either we abandon double aspect PPP, or we show that, within our theoretical framework, change is possible without being caused by a mind. Frankly, the latter turns out to be rather trivial: within MBM, the configuration changes all the time. But change of the configuration simpliciter is not enough. It is less trivial to relate configurations to each other in such a manner that not only individual worlds, but entire, historically consistent world histories can emerge. This is what we will discuss now.

First, we should agree that “brains”, as we usually conceive of them, are but one possible way for objects to embody minds within the framework of double aspect PPP. In principle, there could be many physical systems from which minds of some sort could emerge. We don't need to go into the details about this here. All we need is to agree that mind-bearing physical systems don't necessarily have to resemble brains like ours. To reflect this, I speak of brain-equivalent objects (BEO) rather than “brains”.

What's more, BEOs aren't necessary for the construction of worlds and world histories from configurations and configurational changes. They are only relevant when the phenomenal experiences of mind-bearing beings (such as persons) and their emergence within the framework of MBM are under investigation. Abstracting from such special phenomena, all we need in the more general case are objects that are macroscopic with respect to the level of particles and their configurations – i.e. the fundamental, microscopic level of MBM's ontology – such that they do not necessarily undergo macroscopic change as soon as they undergo microscopic change, but may be affected by microscopic change such that they do undergo macroscopic change. Since BEOs are just a special type of macroscopic objects (MO), what we are going to say about the latter with respect to worlds and world histories carries over to the former. In addition, BEOs will be able to have *experiences*, while MOs, in general, will not, (unless they are BEOs).

Next, we should make the definition of a world a bit more precise. First of all, we need to be clear about what parts of the configuration of the universe are relevant to a MO. As we will see, it turns out that ultimately, this is only the part of the configuration that makes up the MO itself. Whatever happens outside the MO is completely irrelevant to the MO as long as it does not affect the MO, i.e. the part of the configuration that makes up the MO.

If that sounds strange at first, think of the prominent brain in a vat scenario.⁹⁰ Incidentally, the MBM framework allows for scenarios that lie far beyond the outlandishness of this thought experiment. A quick reminder: in the brain in a vat scenario, we imagine a brain that is hooked up to a device, let's say a super computer, that generates neuronal input to the brain such that the brain hallucinates having a body and strolling along a beach, for example, while it actually sits in a vat filled with a liquid that provides for all its physiological needs, including nutrients, heat, and so on. Neuronal output of the brain is fed back to the computer

⁹⁰ See Putnam 1981, 'Brains in a vat'.

and integrated into the simulation, so that the brain doesn't have the impression of living a movie, but instead of actually being able to do things with its body, manipulate the environment, and so on.

In this scenario, some implicit assumptions are made that, strictly speaking, are unnecessary within the MBM framework.

Firstly, we can cut off the connection to the simulation device. Underlying the idea that such a device is required is the assumption that a brain needs some input that is causally connected to some external system in order to experience either something that is not absurd, or even to experience something at all. (Dreams might be an attempt to find a counterexample to the latter, although even when dreaming, we are not completely deprived of sensory input.) But within MBM, what the brain – or BEO, in general – experiences solely depends on either a) the configurational state it is in at the instant it is actualized (Barbour) or b) the succession of configurational states within the world history it lives in (my account), as we will see later. But in any case, the BEO needs no external machine in order to experience a walk at the beach. It just needs to be contained in an actualized a) time capsule or b) world history in which it just happens to have that experience. There is no causal connection to anything external necessary.

Secondly, as an extension of what we have said about the computer, we can get rid of the vat, too. Introducing a vat in which the brain receives everything it physiologically needs to function properly is only required if we make the (implicit) assumption that a brain cannot exist on its own, but must be embedded causally in an environment that first of all allows for its existence and functioning. Under this assumption, an environment too hot or cold, e.g., would already ruin the scenario, for the brain would not function well or even die.

Making this assumption is not necessary in MBM, though. Within each individual, static configuration, there is an essential causal independence between the

configuration of the particles that make up the BEO (or MO) in question and the configuration of the particles that make up all the rest: Whatever exists elsewhere within the configuration containing the MO is irrelevant for the MO's configuration. Thus, in principle, it is irrelevant for the existence of the MO, and, over the course of many actualizations, irrelevant for its experience (if this MO is a BEO capable of experience). In principle, the MO can be the only thing that exists in the universe, and still undergo the configurational change it just happens to undergo. And if this MO is a BEO, it can, accordingly, be the only thing that exists in the universe, and still experience whatever it just happens to experience; it needs no energy, or nutrients, or neural input from the outside. Due to the stochastic micro dynamics of the particles that make up the universal particle configuration, any object composed of a subset of those particles is, in principle, dynamically independent of all the other particles and objects in that configuration. So we can imagine a BEO, e.g. a brain, floating in outer space while having the experience of strolling at the beach. According to the time-capsule approach (a), as soon as this happens the right way, i.e. with the configuration of the brain containing the appropriate snapshots, this brain would indeed have the required experience. Following the world-history approach b), it would take a sequence of changes between configurations for the brain to have that experience, but once the strange thing has happened and the brain has materialized in outer space, there might as well be further actualizations, such that the respective world history emerges. This way or that, this scenario certainly is possible, both in Barbour's account as well as in MBM.

Speaking of world histories, I shall now take the next step towards their construction in order to flesh out the remarks above, in which I appeal to their existence. As has been noted, in order to account for world histories emerging from randomly actualized configurations, we must differentiate between two types of overlap between worlds: overlap simpliciter and overlap proper. Only the latter will prove to be sufficient for world histories to arise, due to the fact that overlap simpliciter does not guarantee a FAPP-continuous evolution on the emergent level. To render it clear just when worlds in MBM overlap properly, and when they don't,

we first have to take a look at MBM's configuration spaces and their inhabitants (i.e. possible configurations), for worlds will be constructed from the latter within the former. The relevant considerations unavoidably will include some digressions into particle metaphysics – particle identity and impenetrability in particular – especially in light of the desideratum of particle number conservation in non-relativistic quantum theories, such as MBM.

IV.7 Configurations and configuration spaces

For a universe of N particles with d degrees of freedom, we can construct a dN -dimensional configuration space \mathbb{R}^{dN} . For physical space with 3 degrees of freedom, this amounts to a $3N$ -dimensional configuration space \mathbb{R}^{3N} .

Let's assume for the moment some sort of principium identitatis indiscernibilium (PII) for physical particles, stating that particles are regarded as identical when they are indistinguishable, and as non-identical when they are distinguishable. For particles to be distinguishable, they need some kind of feature with respect to which they differ from each other. So, for example, an ontology where all particles are featureless, or an ontology where particles aren't featureless, but all instantiate the same features, such that they do not differ despite, must construe them as identical with each other. (For the time being, we are ignoring their spatiotemporal trajectories, which in some cases might be used to individuate them despite their featurelessness, or sameness with respect to their features. We will discuss this later.) Contrariwise, for an ontology where all particles are individuated, for example by their intrinsic properties, haecceity, or a little label (a rigid designator) that is attached to them, there are no identical particles except for the rather trivial case of self-identity. In between these examples lies the case where particles neither are equipped with primitive thisness or the like, nor are all featureless (or all have the same features), but have features – e.g. mass, electrical charge, etc. – in which they differ. Via these features, types of particles can be distinguished, e.g. electrons and positrons, but unless every particle is equipped with a unique

combination of features, or with haecceity, this does not allow for the individuation of all particles there are. Two electrons, say, that do not differ in any of their features, cannot be individuated, although they are numerically different, and even reside at different positions in physical 3-space: just imagine a two-particle universe, with the two particles having identical features. If a demon had exchanged the particles a number of times, putting each at the position of the other and vice versa, without us knowing how often it did that, we would not be able to tell whether the particles in the end are at their original positions or had their positions swapped. There is nothing that could make any difference between the two possibilities. This is what we mean by saying that the particles cannot be individuated.

Depending on the system in question and the particles involved, our representation of their configurations in configuration space may differ. Let's start with the simple case where all particles can be individuated due to them each possessing a unique combination of properties, a label or what have you. In this case, permutations of particles always amount to a different configuration. Hence, all of configuration space is relevant. For an illustration, we assume a toy universe with two individuated particles and one spatial degree of freedom. In this toy universe, the complete two-dimensional plane of our configuration space \mathbb{R}^2 is relevant. Since particles at spatial positions q_1 and q_2 are individuals qua their discerning features, configurations $Q(q_1, q_2)$ and $Q'(q_1, q_2)$ (where particles have changed places, i.e. have exchanged their position values, such that q_2 takes on the value of q_1 and vice versa) are discernible as well, hence distinct. They amount to different facts. Consider the following example (figure 11): One particle resides at $q_1 = -1.5$ and another at $q_2 = 1$. In configuration space, this configuration is represented as the point $Q(-1.5, 1)$ (the blue dot). Now let them change places, such that $q_1 = 1$ and $q_2 = -1.5$, represented by $Q'(1, -1.5)$ (the orange dot).

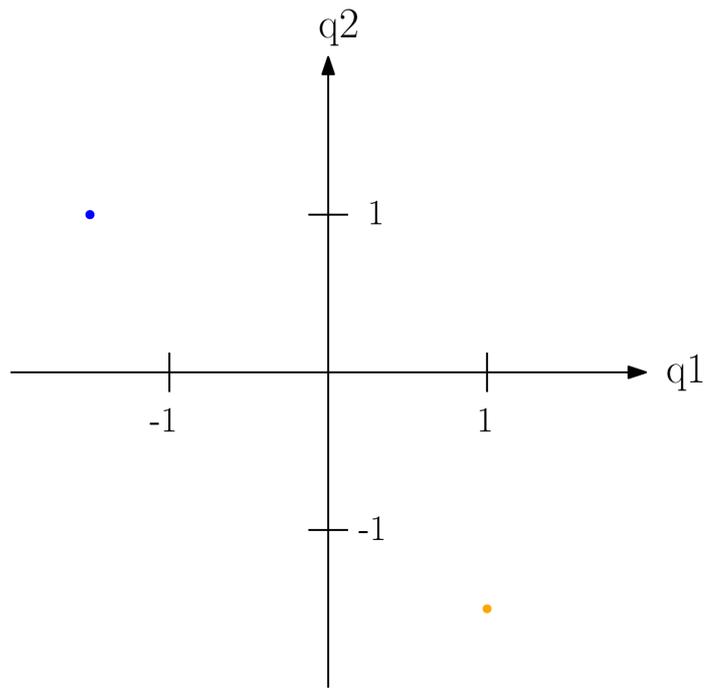


Figure 11: Configuration space \mathbb{R}^2 of two discernible particles with one spatial degree of freedom. The orange and blue dots represent two distinct configurations, despite the fact that in both configurations, the same two spatial locations are occupied.

But what if we are dealing with indiscernible and hence identical particles at q_1 and q_2 ? Q and Q' then should amount to the same configuration, since it doesn't make any difference whether one is at -1.5 and the other at 1 or vice versa. That is to say, mere permutations of identical particles do not individuate configurations. Hence, when representing configurations of identical particles, our configuration space is structurally richer than the possible particle configurations it represents.

Arguably, given identical particles, one should use reduced configuration space to represent meaningfully different possible configurations.⁹¹ To make the formal representation and what it represents match again, we reduce the full configuration space such that the superfluous structure is ignored. For an illustrative example

⁹¹ However, the consequences this has for the WF and its evolution, for the assessment of probabilities, etc., I cannot oversee. In the following discussion of the reduction of configuration space, I therefore only take into account issues of particle identity. Whether, when using a reduced configuration space, the rest of the formalism has to be adapted, and if so, whether that's possible, I leave to the (mathematical) physicists.

(figure 12), we can do this by simply mapping every configuration below the symmetry axis (red) to the other side:

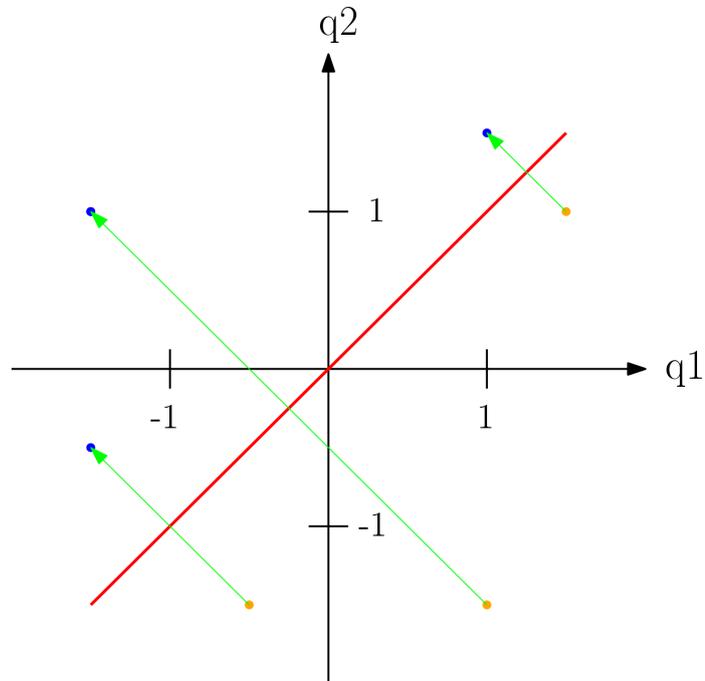


Figure 12: Identical configurations are mapped onto each other in order to reduce superfluous structure.

For higher dimensional configuration spaces, we can do similar rearrangements. If we expand our toy universe by adding a third particle, say, such that all three particles are identical, then there are $3! = 6$ possible permutations that amount to the same configuration. (In general, there are $N!$ permutations for N identical particles, all amounting to the same configuration.) In order to fix that, we get rid of all these permutations except one.

If we do this to all *possible* configurations, i.e. all points in configuration space, we arrive at the so-called reduced configuration space representing all possible configurations of a given system of identical particles that differ meaningfully, i.e. are not mere permutations (Brown et al., 1999, 230): “The ‘reduced’ space \mathbb{R}^{Nd}/S_N [...] is the quotient of \mathbb{R}^{Nd} obtained by the action of the symmetric group S_N (the group of permutations P above)”, the latter simply being defined (ibid., 229) as “an

arbitrary permutation on the set $\{1, 2, \dots, N\}$ with N being the number of particles involved.

IV.7.1 Particle identity, impenetrability, and the desideratum of particle number conservation in MBM

In addition to the mere permutations, it is argued that we also have to remove the diagonals where particles coincide (ibid., 231), i.e. “all points corresponding to two or more particles occupying the same spatial position at the same instant”, resulting in the further reduced configuration space $\mathbb{R}^{Nd}/S_N - \Delta$. As noted by Brown et al. (ibid.), “[this] removal [...] is sometimes justified by considering the particles to be ‘impenetrable’, but of course impenetrability is not a direct consequence of indistinguishability.” Rather, as French and Krause (2006, 40) argue, referring to Newton, “impenetrability is a fundamental assumption of classical physics and we find it explicitly stated in a number of basic texts. Thus Newton characterized the elements of his ‘proto’-theory of classical mechanics in terms of (1) impenetrability, (2) mobility and (3) the ability to excite the senses.” Nevertheless, impenetrability is a metaphysical assumption that might resonate with our intuitions, but needs further justification once we want to make use of it in the construction of our theories about nature. Consequently, French and Krause (ibid., 41) turn the argument around, stating that “for point particles in three dimensions, the regions of coincidence in phase space have zero measure, so impenetrability is a reasonable assumption.” Brown et al. (1999, 231) are calling this line of argument “standard wisdom” when it comes to the construction of reduced configuration spaces, referring to Wu (1984, 2104), who argues that “because we are not always guaranteed that there is a finite probability for two particles to coincide with each other, the so-called diagonal points in \mathbb{R}^{dN} [...] have to be excluded too.”

However, it is questionable whether this argument from the non-finite probability of particle coincidences, and likewise its measure-theoretic analogue brought forward by French and Krause, can convincingly justify the assumption that particles

are impenetrable: belonging to a region of zero measure and hence being assigned a non-finite probability is not to be equated with being impossible. The simple case of a system consisting of only one point-like particle should make that clear. Wherever this particle is located, the region it occupies is of measure zero, unless we assume a coarse-grained cell structure with cells of finite measure for the definition of regions in configuration space, assigning a finite region to the particle's location. Consequently, without imposing such coarse graining, the particle has probability zero of being where it is. Yet, it is there. Clearly, being assigned probability zero or having measure zero are not the same as being impossible. The scope of this argument is not limited to the one-particle case. It holds just the same for two or more particles coinciding, and in general, whenever we are dealing with sets of configurations that have zero measure/are assigned non-finite probability. These sets don't even have to be small, despite their zero measure. The diagonal in question, for example, has the same cardinality as the whole of configuration space.

But then, if these configurations are neither impossible a priori nor of negligible number, why exclude them from configuration space? In a footnote, Wu (1984, fn. 12) points to the fact that, if we didn't exclude the diagonal, we would only be able to obtain Bose-Einstein statistics. This type of argument is different from the ones we have encountered so far. The question of which statistics is adequate for the phenomena to be described is one that can be tackled empirically. If we have to reduce configuration space by the diagonal in order to arrive at a model that captures the phenomena, then we should do so. On these rather solid grounds, then, we may start building a suitable metaphysics for that model, arguing, for example, that the fact that the model, by not allowing for the coincidence of particles, makes correct statistical predictions, should be taken as evidence that the particles in question are impenetrable.

Yet another argument for the subtraction of the diagonal, probably the most important one, goes as follows. Let's say two particles inhabit the same point in space-time. Assume that these particles have no discerning features, such that they

are rendered identical by the above PII. Then, because of them coinciding, they are rendered identical even with respect to their spatiotemporal location, resulting in numerical identity, unless we assume haecceity.⁹² So instead of two particles, we'd end up with only one, either featureless, or, given the original particles had features at all, e.g. mass, with one particle that has double the mass. In any case, our 2-particle system would now consist of only one particle, a situation described by a completely different model (Q , WF , SE) in the first place. This inconsistency could only be resolved by assuming primitive thisness/haecceity, which would allow the preservation of the number of particles despite their coincidence. But if we are assuming haecceity in the case of coincident particles in order to keep N constant, why don't we assume haecceity for particles in general? In fact, I think we would have to, not least because we don't know which particles actually are going to coincide in a future configuration, or did coincide in an earlier one. In standard BM, we do know these things, since there, particles are assumed to travel on trajectories given by the first order GE, which tells us that particles coinciding at one time do so at all times – and conversely, that particles not coinciding at one time never do. Therefore, in standard BM, once we know which particles are coinciding and which aren't, our knowledge about that fact is stable. Not so in MBM, where, in principle, there's nothing to prevent particles from coinciding at one time and not coinciding at others. So, in order to keep N constant, we would need to assume haecceity for all particles. Otherwise it would be possible that two or more non-haecceitistic particles coincide, thereby reducing N . Hence, we really should assume haecceity for particles in general, if we make that assumption at all. But there's a catch: once we assume haecceity for all particles, the reduced configuration space is no longer

⁹² Numerical discernibility of coinciding, identical particles is lost, because they are not even weakly discernible, which would be sufficient for them being numerically discernible. Saunders (2003, 293f.), referring to Black's thought experiment of two iron spheres (Black 1952), provides the irreflexive relation 'x being one mile apart from y' as an example for a weakly discerning feature: "An identity condition may fail even when objects have exactly the same properties and exactly the same relations to all other objects *and* exactly the same relations to each other; [...] Call objects not absolutely or relatively discernible, that satisfy an irreflexive relation, *weakly discernible*; [...] For an example of weakly discernible objects, consider Black's two iron spheres, one mile apart, in an otherwise empty space. [...] The irreflexive relation A is ' \dots one mile apart from \dots '. It is *because* this relationship holds that we may say that there are two [...]"

appropriate to model the N -particle system in question. We'd have to use the unreduced configuration space, because particle permutations amount to different facts about the world, qua the haecceity of particles. This renders the initial reason for the assumption of haecceity absurd. We considered this option in order to be able to keep the diagonal while reducing configuration space. But if we make the assumption, we block the possibility of reduction in general, with or without the subtraction of the diagonal. So once we assume haecceity *in order to* be able to reduce configuration space in a certain way, we can't reduce at all. Therefore, assuming haecceity for that rather formal purpose is a non-starter. Rather, particle haecceity, being an independent, metaphysical assumption, requires strong justification of its own. Whether such justification can be given I shall discuss now.

IV.7.2 Particle haecceity vs. haecceitism

Following Cowling (2016), haecceitism is the position "according to which the world could differ non-qualitatively without differing qualitatively." For an example related to our discussion, consider, in analogy to Black's iron spheres, a universe consisting of two intrinsically identical particles at some spatial distance. This, in the example, is the actual world, and as such, a possible one. Now imagine another possible world where these two particles have swapped places. Haecceitism is the claim that these two worlds exist in the space of possible worlds, and that these possible worlds are different, despite being qualitatively the same.

Based on this characterization, I'd like to distinguish *haecceitism* from *particle haecceity*, the claim that particles *actually* differ in spite of sharing all of their properties, i.e. that the respective particles do have some non-qualitative individuating features to them. Haecceitism with respect to particles, as I am using the notion, merely makes the weaker claim that it is *possible* that they do. At this point, it should be made clear again that we are interested in whether or not we have good reasons to assume particle haecceity, not so much whether particle haecceity, in principle, is possible, i.e. haecceitism. This complicates the issue to

some extent. Chisholm's Paradox, for example, as Cowling (2016) notes, is one of the most prominent arguments brought forward in favour of haecceitism. Unfortunately, though, it is of not much use for our purposes, for it seems to presuppose the individuality of entities in question and their trans-possible-world identity (Chisholm 1967, 3): "The [entity] of this world, we are assuming, is identical with the [entity] of that one" and "moving from one possible world to another, but keeping our fingers, so to speak, on the same two entities [...]" are just two quotes that point towards this presupposition. In more formal presentations of Chisholm's Paradox, this issue shows, appropriately, in the formalism. For example, the two-valued predicate M in the presentation given in Forbes (1984) assigns some set of properties or constituents h_i to some entity α . While h_i changes, thereby representing change of the set of properties or constituents assigned to the entity, α remains the same over the entire operation, the latter being the switching between possible worlds. It is, by assumption, always the same, individuated entity, through all possible worlds. But for our purposes, we must not allow ourselves to make such assumptions, for what is assumed here is exactly the subject matter of the open question we are discussing in this paragraph. Simply assuming individuality and/or trans-world identity would be begging the question.

Other arguments in favour of haecceitism, according to Cowling (2016, 4.1), are mostly of the conceivability type, characterized as follows: "Conceivability arguments for haecceitism have two steps. The first step requires our success in conceiving or imagining certain states of affairs. The second step requires an inference from the relevant conceiving or imagining to the possibility of the states of affairs in question." We are not going to discuss these types of arguments here in detail, for they do not promise to be exceptionally relevant for the special case we are interested in. Generally speaking, when it comes to the question of the haecceity of particles, conceivability type arguments roughly will go as follows: (P1) It is conceivable that two (spatio-temporally identical) particle configurations differ, despite the particles in question being qualitatively identical. (P2) If something is conceivable, it is possible. (C) It is possible that two (spatio-temporally identical)

particle configurations differ, despite the particles in question being qualitatively identical.

Now, regardless of the actual truth-value of P1 and P2, assuming they are true, C is true as well. Nevertheless, the argument has no bearing on our issue at hand. It might well be that configurations differing only with respect to the primitive identity of their particles are conceivable, and it also might well be that, since they are conceivable, that they are possible. But the question we have been trying to answer is different. It is: do particles actually *have* primitive identity? Is their fundamental nature *in fact* haecceitistic? We wanted to be presented with good reasons to assume particle haecceity. What we got instead was the assertion that, given the assumption of some premises, haecceity is possible, i.e., an argument in favour of *haecceitism*, understood as the *possibility of haecceity* – a conviction we shared from the beginning.

We will not ultimately resolve the issues surrounding particle haecceity and general haecceitism here. If the reader feels justified to assume particle haecceity for Bohmian particles, for them, the issue of particle number conservation of coinciding, identical particles in MBM is off the table. Yet, I am unable to find convincing arguments in favour of the haecceity of particles. Accordingly, I have no justification to make use of haecceity in order to solve the problem. So, particle number conservation remains an issue. Maybe we can find a solution by investigating so-called space-time individuality (STI) of particles.

IV.7.3 Individuation of particles qua their space-time trajectories

Earlier, I have mentioned the possibility of particle individuation qua space-time trajectories. Again, I shall make use of an illustrative example. Imagine a universe with all identical particles. At t_1 , we find the particles in configuration Q_1 . At a later time t_2 , we find them in Q_2 . Now, let's assume that Q_1 and Q_2 are identical, in the sense that all points in space that have been occupied by a particle at t_1 are equally

occupied at t_2 . Since particles are identical and thus not discernible through their intrinsic properties, the configurations, too, are indiscernible, despite the particles having changed places. Yet, so the argument goes, there is a way to tell these particles, and a fortiori the configurations, apart: In general, on their way from Q_1 to Q_2 , particles were travelling on different trajectories. Once we know these trajectories, we can regard them as a feature of the particles that renders them discernible. French and Krause (2006, 1.7) reconstruct the conditions under which such a trans-temporal re-identification of particles via their trajectories can be successful. As the first amongst these conditions, they mention spatio-temporal continuity of the trajectory. And some pages later, in their discussion of classical statistical mechanics (ibid., 2.1), they describe the analogue issue that shows there:

“Given that the molecules are all indistinguishable in the sense of possessing the same intrinsic or state-independent properties, some further principle of individuality is required in order to even talk about distinct trajectories in the first place. Boltzmann, more philosophically minded than many of his contemporaries, was completely aware of this issue. In his treatise on the principles of mechanics, the very first axiom of mechanics states that indistinguishable particles which cannot occupy the same point of space at the same time can be individuated by the initial conditions of their trajectories and the continuity of their motion. This, Boltzmann insisted, ‘... enables us to recognize the same material point at different times’.”

This position is what French and Krause call space-time individuality (STI), and, as they note already in the introduction (ibid., 12), “it immediately invites speculation as to the nature of the individuating spatio-temporal background. On a relational view of space and time, the spatio-temporal locations of physical individuals involve relations with other individuals and a possible circularity develops. This may be broken by invoking a privileged set of continuants by reference to which all other things are individuated. But the nature of this privilege must be spelled out and the danger is that STI collapses into some other account of individuality [...].

Alternatively, the circularity could be broken by adopting a substantialist view of space and time according to which spatio-temporal properties and relations of physical things are reduced to the properties of space-time points. Of course, this in turn requires some explication of what it is that confers individuality upon the points of space-time themselves.”

We don't have to delve into these further issues here, for it already should be clear why STI, in our case, can't be utilized for the individuation of particles: there simply are no continuous particle trajectories in MBM, so the necessary condition of STI isn't met. Particle individuation via space-time trajectories, thus, isn't a possible solution to our problem.

So, what are we left with? First, as we have just seen, we can't make use of STI to individuate otherwise identical particles within the MBM framework. Second, while not being strictly impossible, particle haecceity isn't a promising candidate either, for there are no convincing arguments in favour of this assumption.

Yet, there is the formal requirement within the framework of MBM that particle number be conserved, i.e. that N is constant over time, which, without the assumption of haecceity, would be violated once particles do coincide. This requirement is regarded as a postulate of the framework, and cannot simply be dropped, for reasons discussed above. (Recall that two systems with different particle numbers would be described by different models.) The reasonable way out seems to be to assume that particles in MBM indeed cannot coincide, i.e., are impenetrable. Note that this does not contradict the remarks about the relationship between indistinguishability and impenetrability by Brown et al. (1999). Indeed, “impenetrability is not a direct consequence of indistinguishability.” (231) Instead, it is a consequence of 1) the formal desideratum of particle number conservation, on the one hand, and of 2) indistinguishability *plus* 3) the rejection of particle haecceity, on the other. Only if we want all three, we seem to be forced to assume impenetrability: once particles have haecceity or are distinguishable,

impenetrability is not necessary for particle number conservation. But this is true only if *all* particles are distinguishable, given, as we have argued, haecceity is not a reasonable option: as soon as some particles are indistinguishable, conservation of particle number is prone to failure due to possible particle coincidence.

Consequentially, impenetrability must be assumed as soon as two or more particles of a system are indistinguishable, even if we can distinguish particle *types* based on their intrinsic features. To avoid particle coincidences leading to variations in N , we can reduce configuration space not only by the main diagonal, i.e. the configurations where all particles coincide, but also by the secondary diagonals, i.e. the diagonals of the configuration sub-spaces describing the configurations involving two or more, but not all, indiscernible particles. This seems appropriate for modelling a universe that is made up of different, distinguishable particle types, yet allows for indistinguishable particles within those types.

Conclusively, we can say that reduction of configuration space by the diagonals indeed is a formal reflection of the metaphysical assumption of impenetrability. But, as we have seen, at least in MBM, the reasoning behind this assumption is not as trivial as claiming it to be “standard wisdom” or simply referring to our intuitions. Rather, it is a requirement necessary for other desiderata to hold – particle number conservation and no particle haecceity – and only so in the case of indistinguishable particles.

Now that we have gained some clarity about the fundamental space we are operating in, and about the nature of its inhabitants, we finally can approach the ultimate definition of worlds in MBM and see how world histories are construed.

IV.8 What is a world? redux

A particle configuration $Q_i \in \mathbb{R}^{3N}$ is a point given by the positions of the N particles of a given universe. A world $\mathcal{W}_{i(MO)} \in \mathbb{R}^{3N}$ is an area within the configuration space around a configuration Q_i , containing all possible configurations that, for a given

MO within Q_i , satisfy SS with Q_i , i.e., that are sufficiently similar to Q_i , i.e., that are macroscopically indistinguishable from Q_i in the sense that change between them doesn't affect the MO macroscopically, i.e. doesn't make a macroscopic difference to it.⁹³ So let's alter our world membership condition to reflect the MO's relevance:

$$\forall Q_{i(MO)} \exists \mathcal{W}_{i(MO)} \left(\forall Q_j (Q_j \in \mathcal{W}_{i(MO)} \leftrightarrow Q_j SS_{(MO)} Q_{i(MO)}) \right)$$

(world membership_{MBM})

Using this condition, we can define the world $\mathcal{W}_{i(MO)}$ around $Q_{i(MO)}$, containing the MO in question, as the set of all (possible) configurations that, for the MO, are sufficiently similar to $Q_{i(MO)}$ in that their actualization doesn't amount to macroscopic change:

$$\mathcal{W}_{i(MO)} = \{Q_j | Q_j SS_{(MO)} Q_{i(MO)}\}$$

This world trivially contains $Q_{i(MO)}$, since the latter is trivially sufficiently similar to itself.

For defining the world around an *actualized* configuration at a given time, we naturally take for $Q_{i(MO)}$ the configuration actualized at that time.

Note that not necessarily every possible configuration in configuration space has a world around it thus defined. There are only worlds where there are MOs, because it depends on the MO what counts as sufficiently similar (or not) to the actualized configuration. On the other hand, if there is a MO, then there is a world, and if there are several MOs, then there are several worlds. Furthermore, if a configuration contains several MOs, macroscopically being affected by different

⁹³ Of course, it is difficult to even make sense of the notion of macroscopic objects in a toy universe comprising only two particles in one spatial dimension. This toy universe only serves illustrational purposes, to render the underlying principle of the construal of worlds from configurations clear. They are equally applicable to universes with many more particles with more spatial degrees of freedom.

degrees of microscopic change, differently sized worlds emerge from that one configuration.

IV.9 From discrete configurations to continuous world histories via overlapping worlds

We shall now tackle the distinction between overlap simpliciter and overlap proper, and show why only the latter is sufficient for world histories to emerge. Let's start by sharpening our intuitions with respect to overlapping worlds with the help of some illustrations. The crucial feature of worlds as defined above is that they extend through configuration space. Consequently, *worlds can overlap without being identical*.

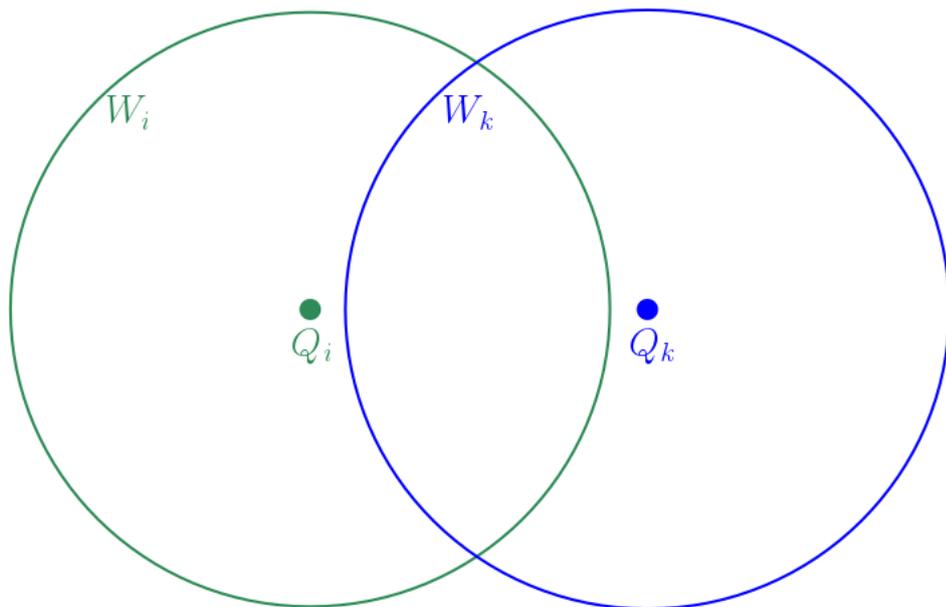


Figure 13: Two worlds overlapping

Figure 13 shows worlds \mathcal{W}_i and \mathcal{W}_k overlapping in configuration space, with $Q_i \notin \mathcal{W}_k$ and $Q_k \notin \mathcal{W}_i$. This is what we call *overlap simpliciter*. From the perspective of the MO in Q_i , the actualization of Q_k would amount to the instantiation of a state differing macroscopically from the one instantiated by Q_i , and vice versa. In this

scenario, no (FAPP-) continuous world history would emerge for the MO. \mathcal{W}_i and \mathcal{W}_k , despite their overlap, in this scenario are two quasi-parallel⁹⁴ worlds.

Only when there is another configuration Q_j actualized within the overlap of \mathcal{W}_i and \mathcal{W}_k can a world history emerge, and FAPP-continuous macroscopic change happen:

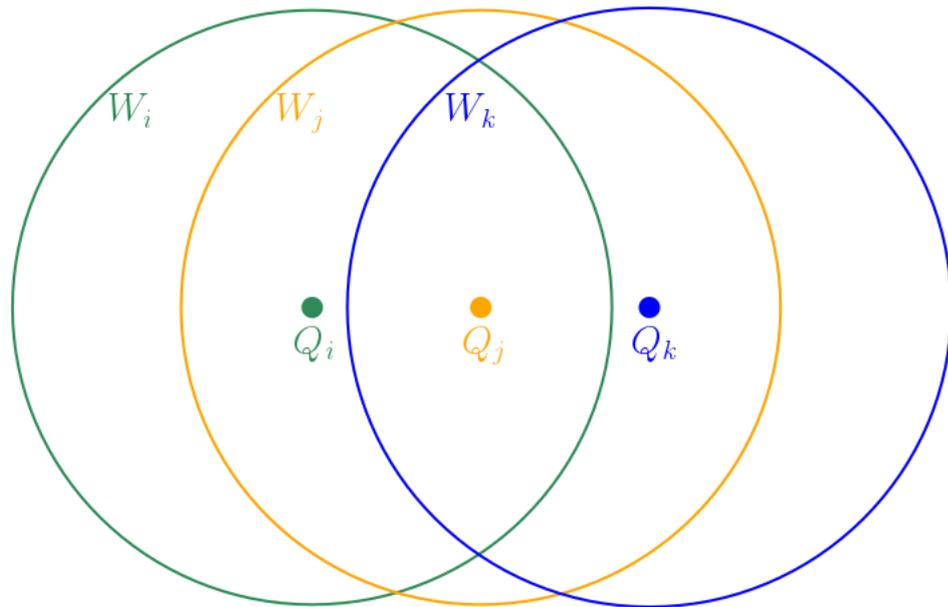


Figure 14: Three worlds, successively overlapping properly, creating a world history

Figure 14 shows the same scenario as figure 13, but now with another configuration Q_j that is actualized in the overlap of \mathcal{W}_i and \mathcal{W}_k , and that includes both Q_i and Q_k in its world \mathcal{W}_j . Now, a FAPP-continuous world history is possible: Let's say we start from Q_i . Q_j lies within \mathcal{W}_i , so if it is actualized, from the perspective of the MO in Q_i , no macroscopic change happens, despite the fact that now, a different configuration Q_j is actualized and hence, a different world \mathcal{W}_j instantiated. Q_i and Q_j satisfy SS, so the MO, from its own macro perspective, persists through the jump between \mathcal{W}_i and \mathcal{W}_j . The same holds for the jump between \mathcal{W}_j and \mathcal{W}_k : now

⁹⁴ „Quasi-parallel“ because in MBM, configurations and worlds are actualized/instantiated sequentially, so worlds, and likewise world histories, do not exist simultaneously, as they do in Everett-style accounts.

residing in Q_j , the MO is not affected by any change when Q_k is actualized. As before, strictly speaking from the fundamental level, the configuration changes; and in accordance with that, different worlds are instantiated. But from the MOs macroscopic perspective, everything remained the same, since Q_j and Q_k satisfy SS.

This is essential. During the configurational jump from Q_i to Q_j , the MO underwent no macroscopic change. Likewise, during the jump from Q_j to Q_k . It macroscopically persists through all jumps, although it doesn't persist on the fundamental level: it quasi- or FAPP-persists. Yet, judging from the perspective of the MO in Q_i , Q_k does not lie within its world \mathcal{W}_i , so overall, with the actualized configuration going from Q_i to Q_k via Q_j , the MO did undergo macroscopic change. Instead of the instantiation of two quasi-parallel worlds \mathcal{W}_i and \mathcal{W}_k , the actualization of Q_j qua instantiating its world \mathcal{W}_j links up \mathcal{W}_i and \mathcal{W}_k in a world history $\{\mathcal{W}_i, \mathcal{W}_j, \mathcal{W}_k\}$, with Q_i, Q_j, Q_k as their respective central configurations. It is only through such chains of actualizations of sufficiently similar configurations, linking up configurations that are not sufficiently similar, that MOs undergo FAPP-continuous, macroscopic change.⁹⁵

Let's make all this a bit more precise. Recall our definition of a world:

$$\mathcal{W}_{i(MO)} = \{Q_j | Q_j SS_{(MO)} Q_i(MO)\}$$

We now need to know when worlds containing MOs overlap. They do when their intersection is non-empty: $\mathcal{O}_{\mathcal{W}_{i(MO)}, \mathcal{W}_{j(MO)}} = \mathcal{W}_{i(MO)} \cap \mathcal{W}_{j(MO)} \neq \emptyset$. Recall our

⁹⁵ Isn't this surprisingly reminiscent of Wittgenstein's account of "family resemblance" in his *Philosophische Untersuchungen*? (Wittgenstein 1953, §67, my translation): "And we are expanding our notion of number, just as we twist fibre on fibre when spinning a thread. And the strength of the thread does not lie in one fibre running through its whole length, but in many fibres overlapping." Of course, instead of „expanding our notion of number“, we are expanding world histories, and the family resemblance we make use of is the relation SS between configurations: We are twisting "fibre on fibre" by having properly overlapping worlds instantiated, and the strength of the resulting thread, i.e. the world history, does not "lie in one fibre running through its whole length" – one continuous trajectory of the configuration – "but in many fibres [(i.e. worlds)] overlapping".

characterization of overlapping worlds in BSM. There, we said that a microstate x_k is in an overlap $\mathcal{O}_{\mathcal{W}_i, j}$ of worlds \mathcal{W}_i and \mathcal{W}_j iff it is in both of these worlds:

$$\forall x_k \forall \mathcal{W}_{i,j} \left(x_k \in \mathcal{O}_{\mathcal{W}_i, j} \leftrightarrow x_k \in \mathcal{W}_i \wedge x_k \in \mathcal{W}_j \right)$$

We can easily adapt this to MBM's configurations:

$$\forall Q_{k(MO)} \forall \mathcal{W}_{i(MO), j(MO)} \left(Q_{k(MO)} \in \mathcal{O}_{\mathcal{W}_{i(MO)}, j(MO)} \leftrightarrow Q_{k(MO)} \in \mathcal{W}_{i(MO)} \wedge Q_{k(MO)} \in \mathcal{W}_{j(MO)} \right)$$

However, this only tells us about overlap simpliciter, i.e. about the necessary and sufficient condition for a micro state to be in an overlap region. For our treatment of BSM, since we were dealing with a continuous trajectory of the micro state through phase space, this was suitable. In such a scenario, the next state instantiated lies arbitrarily close to its predecessor, which guarantees that it is in the latter's world and that thus, the worlds instantiated overlap properly. Yet, for the discontinuous micro dynamics of MBM, we don't have such a guarantee. Hence, overlap simpliciter is not enough. Worlds must overlap properly, i.e. in such a way that their respective central configurations lie within the world around the other central configuration: $Q_i \in \mathcal{W}_j \wedge Q_j \in \mathcal{W}_i$. That is to say, worlds do overlap properly when their respective central configurations lie within their overlap, i.e. within the intersection $\mathcal{O}_{\mathcal{W}_{i(MO)}, j(MO)} = \mathcal{W}_{i(MO)} \cap \mathcal{W}_{j(MO)}$.

Two worlds $\mathcal{W}_{i(MO)}$ and $\mathcal{W}_{j(MO)}$ overlap properly iff the respective central configurations $Q_{i(MO)}$ and $Q_{j(MO)}$ lie within the overlap of these worlds. To express this, we write $\mathfrak{D}_{\mathcal{W}_{i(MO)}, j(MO)}$:

$$\forall Q_{i(MO), j(MO)} \forall \mathcal{W}_{i(MO), j(MO)} \left(\mathfrak{D}_{\mathcal{W}_{i(MO)}, j(MO)} \leftrightarrow Q_{i(MO), j(MO)} \in \mathcal{O}_{\mathcal{W}_{i(MO)}, j(MO)} \right)$$

Having specified proper overlap, we can now construe world histories. Two worlds, despite overlapping properly, don't make a world history, for the simple reason that jumping from one central configuration to the other can, by definition, never amount to macroscopic change for the MO. We only arrive at a world history, e.g. $\mathfrak{S}_{\mathcal{W}_{i,j,k}}$, once we have proper overlap between worlds $\mathcal{W}_i, \mathcal{W}_j$ and proper overlap between worlds $\mathcal{W}_j, \mathcal{W}_k$ but not between $\mathcal{W}_i, \mathcal{W}_k$. There can be more worlds overlapping properly in between, but they don't contribute to the MO undergoing macroscopic change. Also, $\mathcal{W}_i, \mathcal{W}_k$ may overlap simpliciter, but not properly. So, we want to define a world history as a set of worlds such that all neighbouring worlds properly overlap while the first and last world do not:⁹⁶

$$\mathfrak{S}_{\mathcal{W}_{i,\dots,n(MO)}} = \{ \mathcal{W}_{i(MO)}, \mathcal{W}_{j(MO)}, \dots, \mathcal{W}_{m(MO)}, \mathcal{W}_{n(MO)} \mid \mathfrak{D}_{\mathcal{W}_{i(MO)},j(MO)} \wedge \mathfrak{D}_{\mathcal{W}_{j(MO)},\dots} \wedge \dots \wedge \mathfrak{D}_{\mathcal{W}_{\dots,m(MO)}} \wedge \mathfrak{D}_{\mathcal{W}_{m(MO)},n(MO)} \wedge \neg \mathfrak{D}_{\mathcal{W}_{i(MO)},n(MO)} \}$$

Read: the MO's world history $\mathfrak{S}_{\mathcal{W}_{i,\dots,n(MO)}}$ is the set of worlds $\mathcal{W}_{i(MO)}$ through $\mathcal{W}_{n(MO)}$ such that consecutive pairs of worlds overlap properly, but not the first and the last world.

Finally, we have arrived at an account of macroscopically FAPP-continuous change emerging from fundamentally discontinuous actualizations of configurations via the construction of world histories from properly overlapping worlds around individual configurations containing MOs. Note that this FAPP-continuous change throughout a world history in no way necessitates continuous change between configurations. We arrived here starting from the framework of MBM without applying any fundamental changes. In particular, the discontinuous, random actualization of configurations in configuration space, which arguably is the central and discerning feature of MBM, has not been altered.

⁹⁶ Order of the set of worlds in a world history is implicit via the imposed overlap conditions, such that we don't need to construe the set of worlds in a world history as an ordered set explicitly.

A consequence of the emergent nature of macroscopically continuous change, MBM comes with the feature that world histories can be interrupted “externally” without being interrupted “internally”. That is to say: in principle, between two actualizations that belong to one world history as just described, there can lie arbitrarily many actualizations that do not belong to that world history. It could be the case, e.g., that in the time between Q_{j1} and Q_{k1} from $\mathfrak{S}_{\mathcal{W}_{i_1, \dots, n_1(MO)}}$ are actualized, Q_{j2} through Q_{l2} from a different world history $\mathfrak{S}_{\mathcal{W}_{i_2, \dots, n_2(MO)}}$ are actualized. In fact, arbitrarily many configurations, belonging to arbitrarily many different world histories, can be actualized in between the actualization of two consecutive configurations of a given world history. The MO in the latter wouldn’t be affected. Such, I call an external interruption. The world history can be taken up afterwards without noticing, for all individual configurations are static and actualized instantaneously anyway. Internally, world histories are insensitive to external interruption, as long as the conditions of their worlds overlapping properly are met.

How, then, is this proper overlap between the worlds of individual configurations guaranteed, such that a FAPP-continuous world history can emerge? Here comes into play our treatment of the actualization rate of configurations within a world in a finite time interval (cf. II.7): with configurations being actualized at an arbitrarily high rate, the expectation value of configurations belonging to a certain world being instantiated during the same time interval becomes arbitrarily high as well. Apply this to the issue of emergent world histories being FAPP-continuous: given arbitrarily many actualizations within a finite time interval, a configuration being sufficiently similar to a given configuration is practically guaranteed to be actualized within an arbitrarily small amount of time. Once it is actualized, another configuration, being sufficiently similar to the second one, is actualized after an arbitrarily small amount of time, and so on. Always lying within the world of the previous configuration, the consecutive worlds thus instantiated always overlap properly, guaranteeing that world histories continue. It hence is unreasonable to fear that your world history will end just because no configuration within the world

of the presently actualized configuration will be actualized “in time” (as in “punctually”).

IV.9.1 Relative frequency of world histories?

Despite, the relative frequencies with which macroscopically different world histories are instantiated may very well differ in value. Given an arbitrarily high yet finite overall actualization rate, we can, in principle, count the number of actualizations within world histories resembling each other with respect to a certain macroscopic feature during a finite time interval and build a quotient with the overall number of actualizations during that time interval, thus providing a way to “measure” the probabilities of such world histories to be instantiated in terms of frequencies. This is not to say, though, that MBM’s probabilities are defined in terms of relative frequencies. Likewise, this must not be confused with counting branches, or counting world histories. World histories in MBM, like world branches in Everett, cannot be counted. More on that below (IV.10.5). Instead, these relative frequencies are merely *reflecting* the (dispositional) weight of Everettian world branches (cf. Wallace 2012, Ch. 4), providing, so to speak, a specific frequentist way of assessing this weight that is not open to the pure Everettian, who thinks of world branches as being instantiated simultaneously, with their overall combined weight being unity. In short, we are not committing to frequentism with respect to probabilities in the usual sense, but merely use it as a tool to assess dispositional weights of world branches/propensities of world histories to be actualized.

To flesh this out a little, consider an experiment with two possible, mutually exclusive, macroscopic outcomes e and \bar{e} , with probability $P(e) = \frac{2}{3}$ and probability $P(\bar{e}) = 1 - P(e) = \frac{1}{3}$. That is to say, in Everettian terminology, the combined weight of all branches in which e occurs is $\frac{2}{3}$, and the combined weight of all branches in which \bar{e} occurs is $\frac{1}{3}$. For MBM, this translates as follows. If we assume some sufficiently high but finite actualization rate in a given time interval, then the

totality of world histories in which e occurs is instantiated at a relative frequency of about $\frac{2}{3}$ of all actualizations in the time interval considered. Likewise, the totality of world histories in which \bar{e} occurs is instantiated in about $\frac{1}{3}$ of all actualizations in that time interval. This is merely to convey a sense of what it means for world histories to lie in the support of world branches with certain weights, spelled out in terms of relative frequencies of actualizations: world histories in the support of world branches endowed with greater weights are instantiated “more often” than world histories in the support of world branches endowed with lesser weights.

Do these differences in relative frequencies, then, mean that world histories in the support of world branches endowed with very low weights practically cannot emerge? Such fear amounts to a confusion of relative frequency and actualization rate. While the former is approximately stable over all sufficiently high actualization rates – i.e., relative frequencies approximately stay the same, independently of the actualization rate, as long as the latter is sufficiently high – the actualization rate itself can take on arbitrarily high values, such that any world history is actualized arbitrarily often, no matter how low the weight of the respective world branch. And of course, it should be stated explicitly what is already implicit in the assumptions we made about the actualization rate: all this talk of relative frequencies only makes sense given the overall actualization rate is finite. If we assume an infinite overall actualization rate, as we generally do on MBM, we run into the problems familiar from accounts of probabilities in terms of “frequencies” in infinite series of events. For example, assume an infinite overall actualization rate, such that in any finite time interval, infinitely many configurations are actualized. Then, any world history is instantiated infinitely many times in the same time interval. How to calculate, from these infinities, a relative frequency? It might be helpful to imagine the following: If we were able to look at the universe as described by MBM from the outside, but did so only a finite, but sufficiently large number of instants in a given time interval, we'd expect to approximately recover the relative frequencies as above. If, on the other hand, we don't restrict ourselves to looking only a finite number of instants in that time interval, we don't recover relative frequencies.

So much for now about relative actualization frequencies and their role in MBM. We must now come back to world histories, and the ways in which they populate configuration space.

IV.9.2 Quasi-parallel world histories

Above, we have described worlds that do not overlap properly as quasi-parallel worlds. Quasi-parallel worlds can be linked up by world histories. In fact, as argued above, if world histories are to account for macroscopic change, it is crucial that two of their worlds do not overlap properly. Otherwise, no macroscopic change would emerge, since all actualisations within that history would lie within the worlds of all other actualisations. It is for this reason that we included the condition – that two worlds within one world history must not properly overlap – into the definition of a world history. The linking-up of two worlds not properly overlapping is the whole point of world histories, to put it another way.

Just as worlds are said to be quasi-parallel if they do not overlap properly, we can define a condition under which world histories are quasi-parallel. The idea of defining quasi-parallel world histories as those world histories which do not overlap properly suggests itself. But what does it mean for world histories to not overlap properly? Is it necessary that no world from one world history, $\mathfrak{W}_{i_1, \dots, n_1}$, say, properly overlaps with a world from another world history $\mathfrak{W}_{i_2, \dots, n_2}$? Or will it be sufficient if at least one world each does not overlap properly with a world from the other world history, respectively?

To get a grip on this, let's start with the simplest case, in which no world contained in one world history properly overlaps with a world from another world history. First, we need two world histories⁹⁷:

$$\mathfrak{S}_1 = \{ \mathcal{W}_{i_1}, \mathcal{W}_{j_1}, \dots, \mathcal{W}_{m_1}, \mathcal{W}_{n_1} \mid \mathcal{D}_{\mathcal{W}_{i_1, j_1}} \wedge \mathcal{D}_{\mathcal{W}_{j_1, \dots}} \wedge \dots \wedge \mathcal{D}_{\mathcal{W}_{\dots, m_1}} \wedge \mathcal{D}_{\mathcal{W}_{m_1, n_1}} \wedge \neg \mathcal{D}_{\mathcal{W}_{i_1, n_1}} \}$$

and

$$\mathfrak{S}_2 = \{ \mathcal{W}_{i_2}, \mathcal{W}_{j_2}, \dots, \mathcal{W}_{m_2}, \mathcal{W}_{n_2} \mid \mathcal{D}_{\mathcal{W}_{i_2, j_2}} \wedge \mathcal{D}_{\mathcal{W}_{j_2, \dots}} \wedge \dots \wedge \mathcal{D}_{\mathcal{W}_{\dots, m_2}} \wedge \mathcal{D}_{\mathcal{W}_{m_2, n_2}} \wedge \neg \mathcal{D}_{\mathcal{W}_{i_2, n_2}} \}$$

Above, we have specified proper overlap for worlds within a world history. We can likewise specify proper overlap for worlds from different world histories $\mathfrak{S}_1, \mathfrak{S}_2$:

$$\forall Q_{i_1, i_2} \forall \mathcal{W}_{i_1, i_2} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}} \leftrightarrow Q_{i_1, i_2} \in \mathcal{O}_{\mathcal{W}_{i_1, i_2}})$$

We can now differentiate between four basic ways in which world histories relate to each other:

We say that world histories are *entirely quasi-parallel* iff none of their respective worlds overlap properly with any world from the other world history. For this, we write $\mathfrak{S}_1 :: \mathfrak{S}_2$:

$$\mathfrak{S}_1 :: \mathfrak{S}_2 \leftrightarrow \neg \exists \mathcal{W}_{i_1, i_2} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}})$$

(*entirely quasi-parallel*)

Accordingly, we say that world histories are *entirely overlapping* ($\mathfrak{S}_1 \dots \mathfrak{S}_2$) iff each world of the first world history overlaps properly with some world from the other world history *and vice versa*:

⁹⁷ For sake of readability, from now on we will simply write \mathfrak{S}_1 instead of $\mathfrak{S}_{\mathcal{W}_{i_1, \dots, n_1}(MO)}$ etc., and leave out explicit reference to the MO in indices, always keeping in mind that world histories consist of worlds, and worlds only arise from configurations containing MOs.

$$\mathfrak{H}_1 \dots \mathfrak{H}_2 \leftrightarrow \forall \mathcal{W}_{i_1} (\exists \mathcal{W}_{i_2} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}})) \wedge \forall \mathcal{W}_{i_2} (\exists \mathcal{W}_{i_1} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}}))$$

(entirely overlapping)

Word histories are *partially quasi-parallel* ($\mathfrak{H}_1 \dots \mathfrak{H}_2$) iff not every world from the first world history overlaps properly with a world from the other world history or vice versa:

$$\mathfrak{H}_1 \dots \mathfrak{H}_2 \leftrightarrow \neg \forall \mathcal{W}_{i_1} (\exists \mathcal{W}_{i_2} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}})) \vee \neg \forall \mathcal{W}_{i_2} (\exists \mathcal{W}_{i_1} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}}))$$

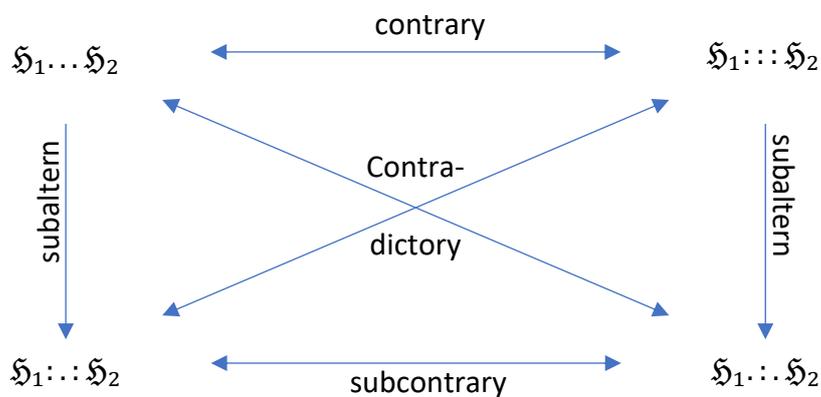
(partially quasi-parallel)

Partially overlapping ($\mathfrak{H}_1 \dots \mathfrak{H}_2$) world histories are such that there are worlds from these different world histories that overlap properly:

$$\mathfrak{H}_1 \dots \mathfrak{H}_2 \leftrightarrow \exists \mathcal{W}_{i_1, i_2} (\mathcal{D}_{\mathcal{W}_{i_1, i_2}})$$

(partially overlapping)

These basic relations between world histories can be arranged in a diagram I call the square of world histories, in analogy to the square of opposition:



The most interesting relation within the square of world histories is probably the subcontrary: two world histories can be partially overlapping and partially quasi-parallel. Stating only the former leaves room for the possibility that the two world histories in question are entirely overlapping. Adding the latter, i.e. world histories being partially quasi-parallel, eliminates this possibility. Likewise, stating only the latter leaves room for the possibility that the world histories are entirely quasi-parallel. But adding the former, i.e. world histories partially overlapping, again, eliminates this possibility. In other words, combining the partially overlapping case with the partially quasi-parallel case allows us to specify a case where two world histories are overlapping in some regions of configuration space, but parallel in others. With this, we have developed the necessary preliminaries to address the topic of splitting and merging world histories.

IV.9.3 World histories: Splitting and merging

Let's take a look at the simplest variant of the case just mentioned, with world histories overlapping in one region of configuration space and being parallel in another. Recall figure 14 above. Figure 15 is the same, with a slight change in indices. To make things simpler (and easier to draw), we imagine that in the region where the world histories are properly overlapping, they do so perfectly, such that, in this region, they are identical even on the fundamental level. While this is not generally so, the condition for overlapping world histories, whether entirely or partially, covers this case.

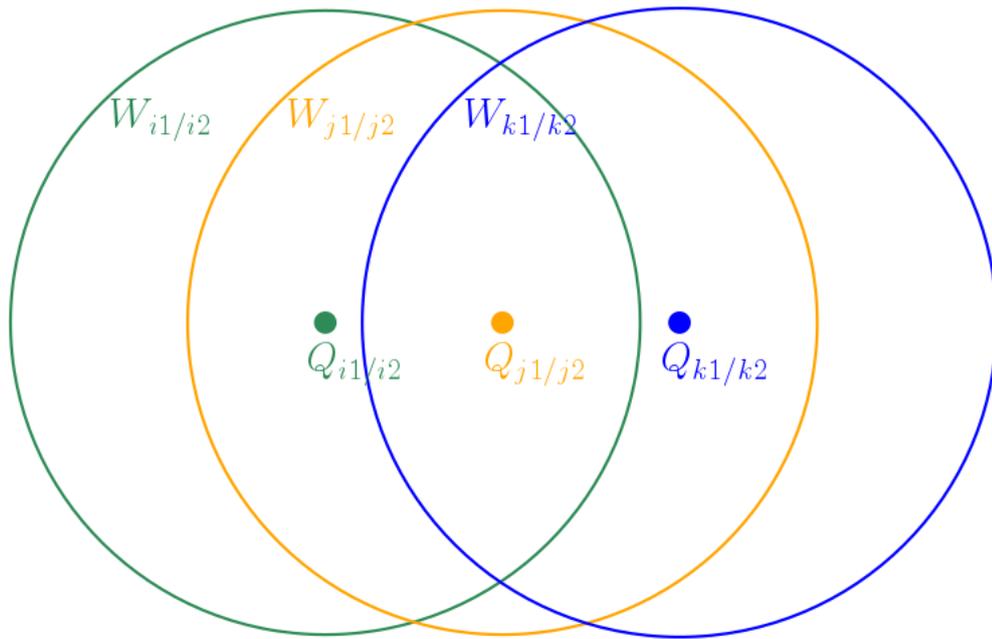


Figure 15: Two (so-far) identical world histories. (Note the indices!)

Now, by adding some more worlds, we can turn this into a picture that contains two world histories that are partially overlapping as well as partially quasi-parallel, i.e. that fall under the conjunction of the mutually subcontrary cases, $\mathfrak{S}_1 :: \mathfrak{S}_2 \wedge \mathfrak{S}_1 :: \mathfrak{S}_2$ (figure 16).

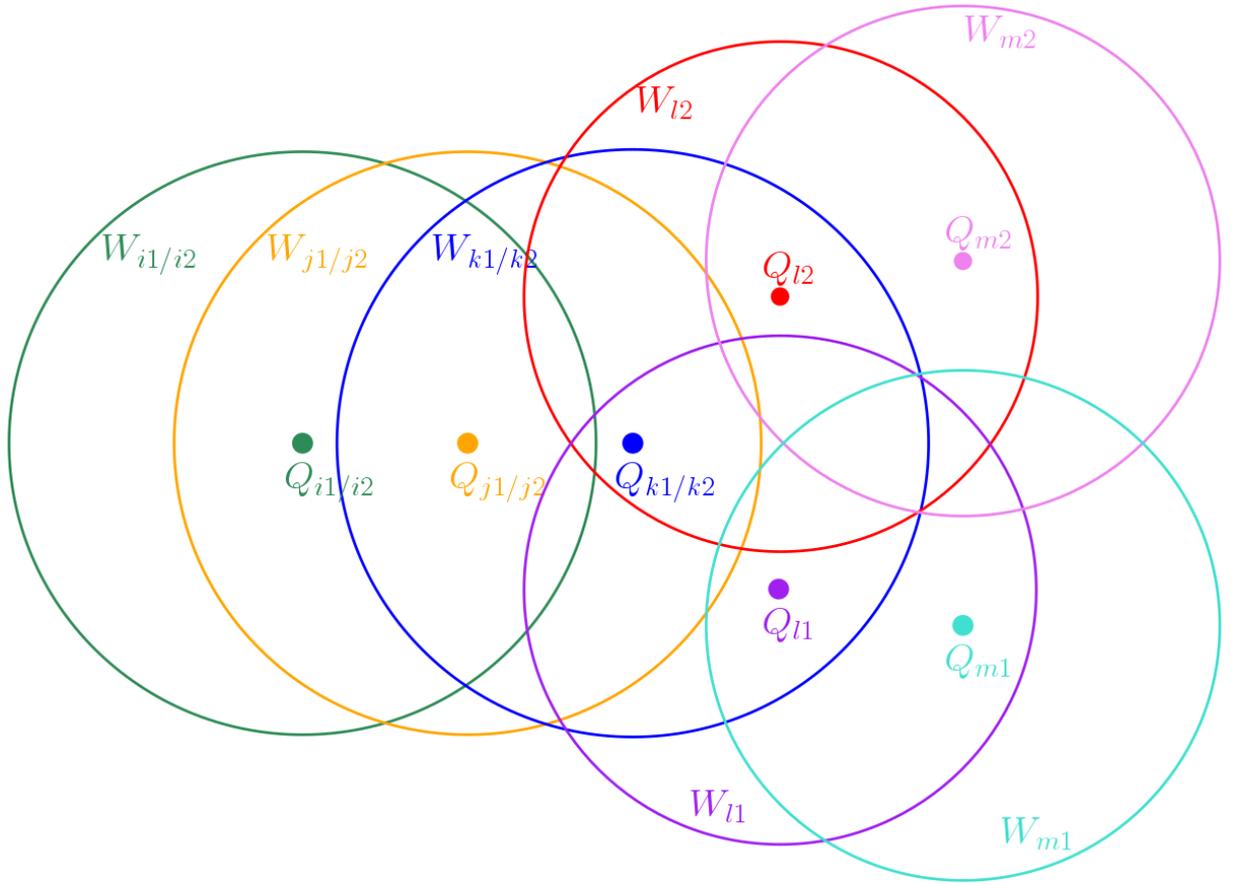


Figure 16: Two world histories, partially overlapping and partially quasi-parallel

Here, Q_{l1} and Q_{l2} both lie within the world of $Q_{k1/k2}$ (and vice versa). As such, they are continuations of the original world history \mathfrak{S}_1 (and of \mathfrak{S}_2 , which, from $Q_{i1/i2}$ to $Q_{k1/k2}$, is identical with \mathfrak{S}_1). But they are not similar enough to lie within each other's worlds. Note that, since worlds \mathcal{W}_{l1} and \mathcal{W}_{l2} do not properly overlap, i.e. neither Q_{l1} nor Q_{l2} lies within the world of the other, world histories \mathfrak{S}_1 and \mathfrak{S}_2 satisfy the condition for partially quasi-parallel world histories. At the same time, they satisfy the condition for partially overlapping world histories, since worlds $\mathcal{W}_{i1/i2}$ through $\mathcal{W}_{k1/k2}$ overlap properly. It is because these two conditions are satisfied that splitting of world histories takes place.

At first sight, there seems to be something peculiar about this picture: read from left to right, it depicts the case of two overlapping world histories (entirely, when looking only at $\mathcal{W}_{i1/i2}$ through $\mathcal{W}_{k1/k2}$), such that they are FAPP one world

history⁹⁸, splitting into two quasi-parallel world histories. Read from right to left, it depicts the case of two quasi-parallel world histories, merging into two overlapping world histories. Which way of reading this picture is the right one? Are these world histories splitting or merging?

The issue is immediately relieved once we remind ourselves that MBM imposes a (temporal) order on its actualizations. To be clear: MBM *prima facie* does not impose a *specific* order on its actualizations. Configurations are actualized at random, i.e. are stochastically independent. Yet, there is only one configuration actualized at a given time, and the next one after that, and so on. Thus, the totality of actualized configurations is ordered, although this order is not predetermined. This ordering carries over to subsets of the totality. Now, focus on configurations that contain MOs and thus instantiate worlds. Since all subsets of the totality are ordered, subsets containing MOs are ordered as well. And since configurations containing MOs instantiate worlds, worlds are ordered, too. If these worlds satisfy the conjunction of overlap conditions as specified above, they form a world history. Therefore, the worlds making up a world history are ordered. We thus know in which direction to read the picture. Assume configurations are actualized in the following order (with possibly other configurations actualized in between): $Q_{i1/i2}$, $Q_{j1/j2}$, $Q_{k1/k2}$, Q_{l1} , Q_{l2} , Q_{m1} , Q_{m2} . From this order, we know that we have to read the picture from left to right. It represents splitting, not merging, of world histories. Note two things: 1) The order just specified is not the only one to instantiate splitting world histories. Any permutation of the last four elements does, if it satisfies $Q_{l1} <_{before} Q_{m1}$ and $Q_{l2} <_{before} Q_{m2}$. That is to say: Within a world history, actualizations of configurations are ordered, but it is irrelevant whether in between these actualizations, configurations from a different world history are actualized or not. Above we have called this external interruption. 2) Ignoring order,

⁹⁸ World histories, in the regions where they entirely overlap, are FAPP one world history, even when they do not overlap as perfectly as depicted in figures 15 and 16. The reason being that entire overlap requires that for every central configuration of one world history, there is another one from the other world history that lies within the former's world. Hence, whenever a central configuration of one of their worlds is actualized, the thus instantiated world overlaps properly with a world of the other world history, such that this instantiation does not amount to macroscopic change for the MO.

there are other possible ways of reading the picture. E.g., starting from \mathcal{W}_{m1} via \mathcal{W}_{l1} to $\mathcal{W}_{k1/k2}$, then splitting into two world histories, $\mathcal{W}_{j1/j2}$ to $\mathcal{W}_{i1/i2}$ and \mathcal{W}_{l2} to \mathcal{W}_{m2} . This and similar possible readings are rendered impossible by the order specified before (and permissible permutations).

A remark about the ordering of actualizations is called for. One might want to know: what imposes this order? I have said that there is one configuration actualized *at a time*. Does this imply that some kind of external, universal, even absolute time is assumed? That is to say, is the order of configurations given by some kind of unambiguous “time stamp”, attached to each configuration as it is actualized, that corresponds to the universal time parameter as it increases? For this presentation of MBM, the answer must be a clear “yes”: indeed, we assume absolute, universal time. The time parameter featuring in the time-dependent SE governing the time evolution of the universal WF suggests itself for such an interpretation. Yet that this assumption is controversial should be clear, too. It is for a reason that Barbour adopts a timeless description of the universe, employing the Wheeler-DeWitt equation. Whether MBM can be altered to be made compatible with such a fundamentally timeless picture of the universe is a question that lies beyond the scope of this thesis. For now, suffice it to say that the assumption of a fundamentally timeless universe has implications on the construal of world histories. With all configurations actualized eternally, we can no longer refer to an external time parameter for their ordering. Hence, the solution applied to the question of splitting and merging world histories above is blocked in such a universe, unless time somehow “re-emerges” on the level of world histories. But for the time being, let’s proceed with our discussion of world histories in MBM, assuming universal time, by applying our framework to the double-slit example familiar from above. This shall prove helpful in making more graspable how world histories emerge from MBM’s discontinuous micro dynamics.

IV.9.4 An application of splitting world histories: The double-slit experiment

Let's look back at figure 16, showing the splitting of world histories. We are now going to use it to model what happens in the double-slit experiment, according to MBM. We start at configuration $Q_{i1/i2}$. That is not to say that the world history really started there. We pick $Q_{i1/i2}$ as a starting point for our considerations, because it is convenient. When this configuration is actualized, it contains all there is in the universe, at that moment in time, from the large scale filaments of galaxies, spanning over hundreds of millions of light years, individual galaxies, to our solar system and earth, including the laboratory and apparatus set up for our experiment, and the particles the behaviour of which we want to investigate. In particular, it also contains the experimenter, carrying out the experiment. At $Q_{i1/i2}$, ever so slightly, they still feel the start button for the apparatus at their fingertip. At $Q_{k1/k2}$, they don't feel that button anymore. And the actualizations in between, $Q_{j1/j2}$ as an example taken from the infinity of others, account for the transition between still feeling the button, in $Q_{i1/i2}$, and not feeling it anymore, in $Q_{k1/k2}$. Let's say that, at the point in time when $Q_{k1/k2}$ is actualized, the particle resides within the slit. To be precise: at that point in time, Q_{k1} , with the particle in the left slit, has some disposition of being actualized, and Q_{k2} , with the particle in the right slit, has some disposition of being actualized. (Remember that these configurations are merely drawn in the figure as the same configuration, while actually being different on the micro level.) With MBM's arbitrarily high actualization rate, configurations sufficiently similar to Q_{k1} and Q_{k2} and thus lying in the worlds \mathcal{W}_{k1} and \mathcal{W}_{k2} are instantiated arbitrarily often: the particle jumps back and forth between the two slits (and other places, too). At least, that is what MBM's micro dynamics tells us. For the experimenter, it makes no difference where the particle is at that moment. It could be somewhere near Andromeda. And, assuming a finite probability for the actualization of that configuration, it will be. As long as these configurations don't amount to a noticeable macroscopic difference for the experimenter, their worlds overlap properly. This apparent curiosity aside, at some point, the experimenter will notice a macroscopic difference: Q_{m1} and Q_{m2} , both being continuations of the original world history, but sufficiently different for the experimenter to be macroscopically distinguishable, instantiate worlds where the

screen is lighting up at different points. Figure 16 conveys the impression that these points should be very close together on the screen, given that the distances between the configurations represented must be miniscule in configuration space. But don't let the limitations of this kind of representation (and of my drawing skills) fool you. You may think of Q_{m1} and Q_{m2} as being as close or far apart as you like, as long as they both instantiate different macroscopic states of the screen, lighting up here or there, and likewise, different macroscopic states of the experimenter. Only with such macroscopic differences splitting of world histories occurs.

Now, think this on into the future. Q_{m1} 's and Q_{m2} 's world histories are continued by the actualization of further configurations. With the experimenter being in Q_{m1} 's world history, say, the experiment goes on, and another point on the screen will light up, leading to branching of world histories again. And the same, of course, for the world history of Q_{m2} . Where the screen lights up in the next event, resulting in the next split of world histories, is just as random as before, "guided" by the universal actualization disposition. But now, the respective world histories already contain records of the first event.⁹⁹ And when these world histories split again, the resulting world histories will contain records of the first two events, and so on, until each resulting world history, at the end of the experiment, contains sufficient records for an interference pattern to emerge on the screen.¹⁰⁰

Two remarks are in order. First, note that the fact that these records are present in future instantiations must not be ascribed to the proper overlap of worlds in the

⁹⁹ Already Bell (1987a, 136, his emphasis) notes that memories and records, telling of different times while being contained in one configuration, can be correlated: "We have only our 'memories' and 'records'. But these memories and records are in fact *present* phenomena. The instantaneous configuration of the x s can include clusters which are markings in notebooks, or in computer memories, or in human memories. These memories can be of the initial conditions in experiments, among other things, and of the results of those experiments. The theory should account for the present correlations between these present phenomena. And in this respect we have seen it to agree with ordinary quantum mechanics, in so far as the latter is unambiguous.

¹⁰⁰ Interestingly, these interference patterns, although being different in each world history branch, can be expected to be roughly the same with respect to their overall structure. They thus grant us a glimpse into the goings-on of the underlying, microscopic level, despite the fact that the world histories we inhabit only emerge on the macro level, and irrespective of which particular world history we inhabit.

respective histories alone. Records, in principle, can vanish in a similar fashion they appear. It is a contingent matter of fact that some records are stable over longer time intervals than others. (The photographic plate that is the screen might vanish after a few centuries; geological formations remain for millions of years; memories of the dreams I have usually last for less than a minute.) Ultimately, stability of records must be attributed to the universal WF and its evolution. However, *if* they are stable, the continuation of world histories via properly overlapping worlds guarantees that they are *reliable*: once a *stable* record is created in a world history, the continuation of this world history will contain this record, too, for as long as it is stable, because worlds that don't contain it are macroscopically distinct and hence not part of this world history. In other words: within a world history, stable records, like those on a photographic plate, do not suddenly vanish. This, although in general, the continuation of world histories via worlds overlapping properly does not impose any kind of (quasi-classical) dynamics onto the universe. This must stem from the dynamics of the universal WF. Proper overlap of worlds *alone* merely guarantees that transitions between worlds are unnoticeable macroscopically, despite a discontinuous micro dynamics, thus explaining the FAPP-continuous evolution of macroscopic processes. In short: world histories take care of *FAPP-continuity* despite a discontinuous micro dynamics of the configuration, while the dynamics of the WF (including decoherence) takes care of *quasi-classicality*.

Second, the example as presented might leave the impression that, on MBM, the experimenter/observer does play a special role with respect to the continuation and splitting of world histories. Yet, this is not the case. They do play a certain role, but only insofar as they are an instance of MOs in general. In particular, it is irrelevant whether the MO in question bears a mind, has conscious experience, etc., or not. We could as well have left out the experimenter from the description, taking the screen itself for the MO. The results with respect to the splitting of world histories would have been decidedly similar, since the observer and the screen are roughly on the same level of macroscopicity: if world histories split for the screen, they split for the experimenter, and vice versa. Then again, viewed from yet another

level, on the order of the solar system, say, no splitting of world histories results from the experiment as described: for a MO on that scale, it literally makes no macroscopic difference whether some screen in some lab on some planet lights up here or there. But of course, things could be arranged such that it does. If we were to “blow up” the result of such an experiment by linking it to an array of nuclear bombs, literally blowing up Mars, for example, iff the first spot on the screen appears in a defined region, and not iff it appears elsewhere, that would indeed make a macroscopic difference for the solar system. Accordingly, in such a scenario, the result of the experiment would lead to splitting of world histories even for the MO that is the solar system.

What this example is supposed to make clear in a more graspable way are the following three points – and the appropriate conclusion – that have been made before: 1. World histories split. 2. Importantly, world histories split, not when some microscopic superposition is present in the state, but when the superposition is magnified to be macroscopically relevant. 3. What is macroscopically relevant depends on the MO in question. In conclusion, the example shows that splitting of world histories is an emergent feature. It naturally results from the MBM framework, i.e. the Schrödinger dynamics for the universal WF, the interpretation of the latter as encoding dispositions for actualizations of configurations, the stochastic dynamics of the actualized configuration in accordance with that interpretation, and the construal of worlds in terms of regions of sufficient similarity with respect to a given configuration, and the MO(s) within it. Accordingly, if there are different MOs for which different processes are macroscopically relevant – like the screen and the solar system in the example – world histories, upon some event, might split or not, because the same event that is macroscopically relevant for one MO might be not for another.

Our treatment of splitting world histories might be suggestive of their countability. I have previously stated that world histories, like Everettian world branches, are not countable. I yet owe an explanation for why that is so.

IV.9.5 Counting world histories

World histories are uncountable. An explanation of this fact is easily given in terms of entirely overlapping and entirely quasi-parallel world histories. Assume two world histories, \mathfrak{H}_1 and \mathfrak{H}_2 , are entirely overlapping, and a third world history \mathfrak{H}_3 is entirely overlapping with \mathfrak{H}_2 . Even if \mathfrak{H}_1 and \mathfrak{H}_3 don't mutually satisfy the condition for entire overlap, and even when they are entirely quasi parallel, they are connected via \mathfrak{H}_2 , similar to the quasi-parallel worlds at the "ends" of an individual world history: \mathfrak{H}_2 , in this scenario, entirely overlaps \mathfrak{H}_1 as well as \mathfrak{H}_3 . Now, is it fair to assume the existence of thusly connected world histories? In this example: is it fair to assume the existence of \mathfrak{H}_2 in between \mathfrak{H}_1 and \mathfrak{H}_3 ? On MBM's infinite actualization rate, it is. It is even fair to assume that there are infinitely many world histories instantiated in between any two quasi-parallel world histories, just as there are infinitely many worlds instantiated in between two quasi-parallel worlds. But, as has been pointed out before, two entirely overlapping world histories can be considered to be, FAPP, one world history. At what point, then, should a world history be considered a different one? How far apart must world histories be in order to count as two world histories, FAPP? The answer is that they must be so far apart that they are entirely quasi-parallel.

Note that both conditions we have made use of rely on worlds as they arise for a given MO. In fact, all conditions for the overlap of world histories do. As specified above, an MO's world depends on the MO itself. Hence, which condition is satisfied – entirely quasi-parallel or entirely overlapping – also depends on the MO. With this dependency made clear, it becomes obvious why world histories cannot be counted, at least not without reference to a specific MO. Returning to the familiar example: for the MO that is the solar system, there are fewer world histories than for the MO that is me. Just which MO we should select as a reference point to count world histories is a question that has no unique answer. Rather, it is a matter of context. Ultimately, world histories, like the worlds that constitute them, are an

emergent feature – they only exist FAPP, from the perspective of a MO.

Fundamentally, all there is are random actualizations of configurations, following an actualization disposition encoded in the universal WF. Despite, or rather: because of this, world histories are suitable to explain the manifest image of FAPP-continuous, macroscopic processes emerging from the scientific image of a stochastic, discontinuous micro dynamics.

IV.10 MBM is compatible with double aspect PPP

As argued above, in order to be compatible with a more than minimal variant of PPP, in particular in order to be compatible with double aspect PPP, MBM must allow for macroscopic processes from which what we would call a mind capable of experiences can emerge. Individual configurations, even if they are time capsules, didn't seem to be enough, for the processes required as the physical aspect of mental phenomena, i.e. brain processes, cannot arise from within individual, static configurations. Employing FAPP-continuous world histories, we are able to account for experiences, including experiences of continuous macroscopic change, despite a fundamentally random, discontinuous micro dynamics. Our brains are MOs undergoing macroscopic change within world histories. For the latter, we require pairs of consecutive configurations to be sufficiently similar as to not affect the MO macroscopically by transitioning from one to the other. At the same time, because of similarity being a nontransitive relation, a series of actualizations of pairwise sufficiently similar configurations can have configurations at its far ends that are not sufficiently similar to each other, and thus amount to macroscopic change of the MO. If the MO is a brain, this allows for brain processes to happen, from which an experiencing mind can emerge, as the psychological aspect of the physical process. It is probably worth noting again that we don't need a mind already in place for world histories to form. In particular, it is not necessary that there is a mind (or a god, even) comparing configurations in order to see whether they are sufficiently similar or not. Sufficient similarity merely requires that two actualisations do not amount to macroscopic change of an MO. Whether they do or not, as argued,

depends only on the MO in question, not on a mind observing it, or the MO itself having a mind. It can undergo change, even in an otherwise completely empty universe, deprived of any mental aspect whatsoever. For example, the screen or photographic plate in the double-slit experiment is an MO that undergoes macroscopic change during the experiment, without itself having a mind, and irrespective of whether an experimenter – a special, experiencing type of MO – is watching or not. While undergoing change, according to the best theories we have, seems to be a necessary requirement for a mind to emerge, neither having a mind nor being observed by a mind is a necessary requirement for undergoing change. We are thus not in need of a mind for world histories to form – MBM indeed seems compatible with double aspect PPP.

V. Conclusion and Outlook

In this thesis, I have developed and presented an account of the emergence of FAPP-continuous, quasi-classical world histories from the stochastic, discontinuous micro dynamics of MBM.

For an introduction to MBM, the transition from BM to MBM was described, and their empirical equivalence established. I presented the conditions under which MBM can be classified as a primitive ontology theory, regarded as crucial for the acceptance of a theory as Bohmian by many proponents of BM. While rendering MBM compatible with the PO approach, it must not fall prey to the problem of communication, arising from the two-space reading of BM. After spelling out the desiderata of the PO approach regarding the fundamentality of the primitives of a theory, I argued that to achieve compatibility with the PO approach while avoiding the problem of communication, MBM's wave function must be interpreted nomologically. That, however, is not to say that I commit myself to this interpretation. I merely argued that MBM can be classified as a PO theory, and thus as a Bohmian theory, if wished: given a suitable reading of the wave function, MBM

satisfies the desiderata of the PO approach. Thus, Bohmians should abstain from rejecting MBM for not possibly being Bohmian.

I have discussed the issue of temporal solipsism, identified by Bell as a serious problem for his “Everett (?) theory” – the historic predecessor of MBM. I argued that Barbour’s time-capsule approach, adopted for MBM, does not provide a satisfactory answer to Bell’s worry. In particular, his adopting a more than minimal psychophysical parallelism between brain processes and experience of macroscopic change, while being a reasonable assumption in light of our current best neuroscience, is problematic for theories relying solely on time capsules, understood as highly structured, individual, static configurations containing consistent records. As a solution, I introduced worlds, and world histories, as key concepts in providing a link between MBM’s discontinuous micro dynamics and FAPP-continuous macroscopic phenomena. Worlds, other than time capsules, are coarse-grained regions in configuration space, defined as sets of possible configurations satisfying a nontransitive relation of sufficient similarity with respect to a given configuration containing macroscopic objects. Hence, worlds may overlap. They thus allow for FAPP-continuous, quasi-classical world histories to form.

To first introduce the notion of overlapping worlds, I discussed their applicability to the case of Boltzmannian statistical mechanics (BSM). There, for finite available phase space volumes, partitioned into non-overlapping macro regions of finite, non-zero volume, the discontinuity problem of BSM arises: BSM invariably models change of macro values in such scenarios as jump discontinuities, at odds both with the expectation that macro values should be able to change continuously – stemming from the desideratum that BSM be able to model thermodynamically described target systems – and with the continuous evolution of the micro state. I have argued that overlapping worlds are a reasonable option for replacing the disjoint macro regions of phase space. This move solves the discontinuity problem.

Having thus introduced and developed the notion of overlapping worlds, and demonstrated its applicability to the case of BSM, I adapted the concept for the case of MBM, where a similar problem arises: FAPP-continuous, quasi-classical, consistent histories (as we seem to live in one) must emerge from MBM's discontinuous, fundamentally random micro dynamics in configuration space, thereby avoiding Bell's charge of temporal solipsism. A special requirement for overlapping worlds had to be fulfilled in order for world histories to form: worlds must overlap properly. I argued that MBM's indefinite actualization rate guarantees that they do. Along the way, issues pertaining to the microscopic-macroscopic distinction, reduction of configuration space, particle identity, impenetrability and haecceity in light of the desideratum of particle number conservation in non-relativistic quantum theories, etc., have been examined. I provided a detailed explanation of how overlapping worlds in MBM form world histories, thereby linking up macroscopically distinct worlds. Thus, the problem of temporal solipsism was resolved.

My hope was that by spelling out in some detail how world histories are formed, MBM might be rendered more acceptable to those having difficulties imagining how the manifest image of the universe we seem to inhabit can possibly emerge from the peculiar scientific image of the universe MBM describes. While working on MBM for some time now, and discussing it with friends and colleagues, I became increasingly convinced that I should at least attempt to make this emergence intelligible. I hope I have succeeded, if only to some extent. Certainly, there are more questions to ask, and to answer, and to answer more extensively and in greater detail, than I could in this thesis. MBM is a new theory, despite descending from a decade old ancestor: Bell's Everett (?) theory. Possibly it was because its original inventor – or discoverer, if you prefer – had dismissed his own theory immediately, and because he was such an eminent figure, that it never really took off. Alas, Bell's criticism was based on a misguided identification of worlds with individual configurations. Had he not made this mistake, MBM by now could have been "out there", as a viable alternative, for decades, possibly developed way

further than the juvenile stage it finds itself in as of today. However, Bell did not recognize that worlds can – and should – be construed as coarse-grained, and so, for the time being, MBM is still a work in progress, with plenty of topics to be taken up in possible follow-up projects. For an example, there are interesting questions to engage with concerning the personal identity of inhabitants of (splitting) world histories: can persons inhabiting remote parts of world histories reasonably be identified with one another? What about persons in parallel world histories? If the latter are a result from splitting, are there now multiple persons where there was one, before the split? Or have there always been several, already before splitting? What are the differences, if any, between simultaneously realized, Everettian world branches, and MBM's world histories, with respect to personal identity? I think such and related questions might be of interest even to a broader audience, not particularly concerned with MBM in itself. Likewise, issues pertaining to the relationship between world-history internal time and external, universal time may be worth investigating. After all, world histories may be externally interrupted without being internally interrupted, as long as their worlds overlap properly. The fact that, on MBM, there is no necessary link between fundamental, universal time and perceived time (except for ordering actualizations) might inform more general debates in the philosophy of time broadly construed. I would like to pick up these and similar topics in future projects, and can only invite the interested reader to join me.

References

Albert, D. Z. (1996). Elementary Quantum Metaphysics. In: Cushing, J. T., Fine, A., & Goldstein, S. (Eds.). *Bohmian mechanics and quantum theory: an appraisal*. Springer Science+Business Media.

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

Allori, V. (2015). Primitive ontology in a nutshell. *International Journal of Quantum Foundations*, 1(3), 107–122.

Allori, V., Goldstein, S., Tumulka, R., & Zanghì, N. (2008). On the Common Structure of Bohmian Mechanics and the Ghirardi-Rimini-Weber Theory. *The British Journal for the Philosophy of Science*, 59(3), 353–389.

Bacciagaluppi, G. (2016). The Role of Decoherence in Quantum Mechanics. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition). <https://plato.stanford.edu/archives/fall2016/entries/qm-decoherence/>

Barbour, J. (1994). The emergence of time and its arrow from timelessness. In: Halliwell, J. J., Pérez-Mercader, J. & Zurek, W. H. (Eds.), *Physical Origins of Time Asymmetry*. Cambridge University Press. 405–414.

Barbour, J. (1999). *The End of Time*. Oxford University Press.

Barbour, J. (2009). The Nature of Time. [arXiv:0903.3489v1](https://arxiv.org/abs/0903.3489v1) [gr-qc]

Barbour, J. (manuscript). *The Janus Point*.

Barnes, L. A. (2012). The Fine-Tuning of the Universe for Intelligent Life. *Publications of the Astronomical Society of Australia*, 29(4), 529–564.

<http://doi.org/10.1071/AS12015>

Barrett, J. A. (1996). Empirical adequacy and the availability of reliable records in quantum mechanics. *Philosophy of Science*, 63(1), 49–64.

<http://doi.org/10.1086/289893>

Beisbart, C., & Hartmann, S. (Eds.) (2011). *Probabilities in Physics*. Oxford University Press.

Bell, J. S. (1987a). Quantum mechanics for cosmologists. In: Bell, J. S. (1987b). *Speakable and Unsayable in Quantum Mechanics*. Cambridge University Press.

Bell, J. S. (1987b). *Speakable and Unsayable in Quantum Mechanics*. Cambridge University Press.

Ben-Menahem, Y., & Hemmo, M. (Eds.) (2012). *Probability in Physics*. Springer Science & Business Media.

Black, M. (1952). The identity of indiscernibles. *Mind*, 61(242), 153–164.

Bohm, D. (1952a). A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. I. *Physical Review*, 85(2), 166–179.

<http://doi.org/10.1103/PhysRev.85.166>

Bohm, D. (1952b). A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. II. *Physical Review*, 85(2), 180–193.

<http://doi.org/10.1103/PhysRev.85.180>

Bohm, D. & Hiley, B. J. (1993). *The Undivided Universe: An Ontological Interpretation of Quantum Theory*. Routledge.

Bohr, N. (1949). Discussions with Einstein on Epistemological Problems in Atomic Physics. In: Schilpp, P. A. (Ed.) (1949). *Albert Einstein: Philosopher-Scientist*. Tudor.

Brading, K., and Castellani, E. (Eds.), *Symmetries in Physics*. Cambridge University Press.

Brown, H. R. (2010). Reply to Valentini, 'De Broglie–Bohm Pilot-Wave Theory: Many Worlds in Denial?' In: Saunders, S., Barrett, J., Kent, A., & Wallace, D. (Eds.). (2010). *Many Worlds?* Oxford University Press, 510-517.

Brown, H. R., Sjöqvist, E., & Bacciagaluppi, G. (1999). Remarks on identical particles in de Broglie-Bohm theory. *Physics Letters A*, 251(4), 229–235.

[http://doi.org/10.1016/s0375-9601\(98\)00907-4](http://doi.org/10.1016/s0375-9601(98)00907-4)

Brown, H. R., & Wallace, D. (2005). Solving the Measurement Problem: De Broglie-Bohm Loses Out to Everett. *Foundations of Physics*, 35(4), 517–540.

<http://doi.org/10.1007/s10701-004-2009-3>

Butterfield, J. (2002). Review: Julian Barbour, *The End of Time*. *The British Journal for the Philosophy of Science*, 53(2), 289–330.

Callender, C. (1999). Reducing Thermodynamics to Statistical Mechanics: The Case of Entropy. *The Journal of Philosophy*, 96(7), 348–373.

Callender, C. (2004). There is no Puzzle about the Low-Entropy Past. In: Hitchcock, C. (Ed.) (2004). *Contemporary Debates in Philosophy of Science*. Wiley-Blackwell.

Callender, C. (2007). The emergence and interpretation of probability in Bohmian mechanics. *Studies in History and Philosophy of Modern Physics*, 38(2), 351–370.
<http://doi.org/10.1016/j.shpsb.2006.08.004>

Carroll, S. M. (2017). Why Boltzmann Brains Are Bad. [arXiv:1702.00850v1](https://arxiv.org/abs/1702.00850v1) [hep-th]

Chattaraj, P. K. (Ed.). (2011). *Quantum Trajectories*. CRC Press.

Chen, E. K. (2019). Realism about the wave function. *Philosophy Compass*, 14(7).
<http://doi.org/10.1111/phc3.12611>

Chisholm, R. M. (1967). Identity through Possible Worlds: Some Questions. *Noûs*, 1(1), 1.

Clark, P., & Read, S. (1984). Hypertasks. *Synthese*, 61(3), 387–390.
<http://doi.org/10.1007/BF00485061>

Cowling, S. (2016). Haecceitism. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition).
<https://plato.stanford.edu/archives/fall2016/entries/haecceitism/>

Cushing, J. T. (1994). *Quantum mechanics : historical contingency and the Copenhagen hegemony*. The University of Chicago Press.

Cushing, J. T., Fine, A. & Goldstein, S. (Eds.). (1996). *Bohmian mechanics and quantum theory: an appraisal*. Springer Science+Business Media.

Dürr, P., & Ehmman, A. (ms). *On the Physics and Metaphysics of Minimal Bohmian Mechanics*. (under review)

Dürr, D., Goldstein, S. & Zanghì, N. (1992). Quantum Equilibrium and the Origin of Absolute Uncertainty. *Journal of Statistical Physics*, Volume 67, Issue 5-6, 843-907. <https://doi.org/10.1007/BF01049004>

Dürr, D., & Teufel, S. (2009). *Bohmian Mechanics*. Springer Science & Business Media.

Dürr, D., Goldstein, S., & Zanghì, N. (2012). *Quantum Physics Without Quantum Philosophy*. Springer Science & Business Media.

Esfeld, M. (2014). Quantum Humeanism, or: Physicalism without Properties. *The Philosophical Quarterly*, 64(256), 453–470. <http://doi.org/10.1093/pq/pqu030>

Esfeld, M. (2018). Collapse or No Collapse? What Is the Best Ontology of Quantum Mechanics in the Primitive Ontology Framework? In: Gao, S. (Ed.). *Collapse of the Wave Function*, Cambridge University Press.

Forbes, G. (1984). Two solutions to Chisholm's paradox. *Philosophical Studies*, 46(2), 171–187.

Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, Neue Folge, Band 100/1, 25-50. Pfeffer Leipzig.

French, S., & Krause, D. (2006). *Identity in Physics: A Historical, Philosophical, and Formal Analysis*. Oxford University Press. <http://doi.org/10.1093/0199278245.001.0001>

Friebe, C., Kuhlmann, M., Lyre, H., Näger, P., Passon, O. & Stöckler, M. (2015). *Philosophie der Quantenphysik*. Berlin, Heidelberg: Springer-Verlag. <http://doi.org/10.1007/978-3-642-37790-7>

Friederich, S. (2018). Fine-Tuning. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition).

<https://plato.stanford.edu/archives/win2018/entries/fine-tuning/>

Frigg, R. (2008). A Field Guide to Recent Work on the Foundations of Statistical Mechanics. In: Rickles, D. (Ed.) (2008). *The Ashgate Companion to Contemporary Philosophy of Physics*. Ashgate.

Frigg, R., Berkovitz, J., & Kronz, F. (2016). The Ergodic Hierarchy. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition).

<https://plato.stanford.edu/archives/sum2016/entries/ergodic-hierarchy>

Frigg, R. & Hartmann, S. (2020). Models in Science. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2020 Edition).

<https://plato.stanford.edu/archives/spr2020/entries/models-science/>

Frigg, R. & Werndl, C. (2011). Entropy: A Guide for the Perplexed. In: Beisbart, C., & Hartmann, S. (Eds.) (2011). *Probabilities in Physics*. Oxford University Press. 115-

142. <http://doi.org/10.1093/acprof:oso/9780199577439.001.0001>

Frigg, R., & Werndl, C. (2012). A New Approach to the Approach to Equilibrium. In: Ben-Menahem, Y. & Hemmo, M. (Eds.) (2012). *Probability in Physics*. Springer Science & Business Media. 99-113.

Frigg, R., & Werndl, C. (2019). Statistical Mechanics: A Tale of Two Theories. *The Monist*, 102(4), 424–438. <http://doi.org/10.1093/monist/onz018>

Gao, S. (Ed.) (2018). *Collapse of the Wave Function*, Cambridge University Press.

Greenberger, D., Hentschel, K., & Weinert, F. (Eds.) (2009). *Compendium of Quantum Physics*. Springer Science & Business Media.

Greenberger, D., Reiter, W. L., & Zeilinger, A. (Eds.) (1999). *Epistemological and Experimental Perspectives on Quantum Physics*. Springer Science & Business Media.

Goldstein, S. (2012). Typicality and Notions of Probability in Physics. In: Ben-Menahem, Y., & Hemmo, M. (Eds.) (2012). *Probability in Physics*. Springer Science & Business Media.

Goldstein, S. (2017). Bohmian Mechanics. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition).

<https://plato.stanford.edu/archives/sum2017/entries/qm-bohm/>

Goldstein, S., Huse, D. A., Lebowitz, J. L., & Tumulka, R. (2017). Macroscopic and microscopic thermal equilibrium. *Annalen der Physik*, 529, 1600301.

<https://doi.org/10.1002/andp.201600301>

Goldstein, S., Lebowitz, J. L., Tumulka, R., & Zanghì, N. (2019). Gibbs and Boltzmann Entropy in Classical and Quantum Mechanics. *arXiv.org*.

<https://arxiv.org/abs/1903.11870v2>

Goldstein, S., Tumulka, R. & Zanghì, N. (2011). Bohmian Trajectories as the Foundation of Quantum Mechanics. In: Chattaraj, P. K. (Ed.). (2011). *Quantum Trajectories*. CRC Press.

Hacking, I. (1987). The Inverse Gambler's Fallacy: the Argument from Design. The Anthropic Principle Applied to Wheeler Universes. *Mind*, XCVI(383), 331–340.

Halliwell, J. J., Pérez-Mercader, J. & Zurek, W. H. (Eds.), *Physical Origins of Time Asymmetry*. Cambridge University Press.

Heidelberger, M. (2003). The mind-body problem in the origin of logical empiricism: Herbert Feigl and psychophysical parallelism. In: Parrini, P., Salmon, W. C. & Salmon, M. H. (Eds.) (2003). *Logical Empiricism: Historical and Contemporary Perspectives*, University of Pittsburgh Press, 233-262.

Hemmo, M., & Shenker, O. R. (2012). *The Road to Maxwell's Demon*. Cambridge University Press.

Hiley, B. J. (1999). Active Information and Teleportation. In: Greenberger, D., Reiter, W. L., & Zeilinger, A. (Eds.). *Epistemological and Experimental Perspectives on Quantum Physics*. Springer Science & Business Media.

Hiley, B. J. (2009). Bohm Interpretation of Quantum Mechanics. In: Greenberger, D., Hentschel, K., & Weinert, F. (Eds.). *Compendium of Quantum Physics*. Springer Science & Business Media.

Hitchcock, C. (2004). *Contemporary Debates in Philosophy of Science*. Wiley-Blackwell.

Jaeger, G. (2014). What in the (quantum) world is macroscopic? *American Journal of Physics*, 82(9), 896–905. <http://doi.org/10.1119/1.4878358>

Kim, J. (1993). *Supervenience and Mind*. Cambridge University Press.

Maudlin, T. (1995). Three measurement problems. *Topoi*, 14(1), 7–15. <http://doi.org/10.1007/BF00763473>

McLaughlin, B. & Bennett, K. (2018). Supervenience. In: Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition). <https://plato.stanford.edu/archives/win2018/entries/supervenience/>

Ney, A. (2012). The Status of our Ordinary Three Dimensions in a Quantum Universe. *Noûs*, 46(3), 525–560. <http://doi.org/10.1111/j.1468-0068.2010.00797.x>

Norsen, T. (2018). On the Explanation of Born-Rule Statistics in the de Broglie-Bohm Pilot-Wave Theory. *Entropy*, 20(6), 422–26. <http://doi.org/10.3390/e20060422>

Norton, J. D. (2012). Approximation and Idealization: Why the Difference Matters. *Philosophy of Science*, 79(2), 207–232. <http://doi.org/10.1086/664746>

Norton, J. D. (2015). You are not a Boltzmann Brain. Retrieved January 25, 2018, from https://www.pitt.edu/~jdnorton/Goodies/Boltzmann_Brain/Boltzmann_Brain.html

Oppenheimer, D. M., & Monin, B. (2009). The retrospective gambler's fallacy: Unlikely events, constructing the past, and multiple universes. *Judgement and Decision Making*, 4(5), 326–334.

Parrini, P., Salmon, W. C. & Salmon, M. H. (Eds.) (2003). *Logical Empiricism: Historical and Contemporary Perspectives*, University of Pittsburgh Press, 233–262.

Passon, O. (2010). *Bohmsche Mechanik*. Europa-Lehrmittel.

Putnam, H. (1981). *Reason, Truth and History*. Cambridge University Press.

Reichenbach, H. (1956). *The Direction of Time*. Dover.

Reif, F. (1965). *Fundamentals of statistical and thermal physics*. McGraw-Hill Book Company.

Rickles, D. (Ed.). (2008). *The Ashgate Companion to Contemporary Philosophy of Physics*. Ashgate.

Sacks, O. (2017). *The River of Consciousness*. Picador.

Salmon, M. H. (Eds.), *Logical Empiricism: Historical and Contemporary Perspectives*, University of Pittsburgh Press.

Saunders, S. (2003). Physics and Leibniz's Principles. In: Brading, K., and Castellani, E. (Eds.), *Symmetries in Physics*. Cambridge University Press.

Saunders, S., Barrett, J., Kent, A., & Wallace, D. (Eds.). (2010). *Many Worlds?* Oxford University Press.

Schilpp, P. A. (Ed.). (1949). *Albert Einstein: Philosopher-Scientist*. Tudor.

Schlosshauer, M. (2007). *Decoherence and the Quantum-To-Classical Transition*. Springer.

Schlosshauer, M. (2014). The quantum-to-classical transition and decoherence. <https://arxiv.org/abs/1404.2635v2>

Schlosshauer, M., & Camilleri, K. (2011). What classicality? Decoherence and Bohr's classical concepts. *Advances in Quantum Theory: Proceedings of the International Conference on Advances in Quantum Theory*, AIP, Vol. 1327, pp. 26–35.

<http://doi.org/10.1063/1.3567426>

Stone, J. (2018). Belgium and the Netherlands swap land to change their national borders. <https://www.independent.co.uk/news/world/europe/belgium-netherlands-national-border-change-meuse-vise-eijsden-maastricht-a8141166.html>

Last accessed September 21, 2019.

- Suárez, M. (2015). Bohmian dispositions. *Synthese*, 192(10), 3203–3228.
<http://doi.org/10.1007/s11229-015-0741-1>
- Tahko, T. E. (2018). Fundamentality. In: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition).
<https://plato.stanford.edu/archives/fall2018/entries/fundamentality/>
- Tao, T. (2011). *An Introduction to Measure Theory*. American Mathematical Society.
- Valentini, A. (1991). Signal-locality, uncertainty, and the subquantum H-theorem. I. *Physics Letters A*, 156(1,2), 5–11. [http://doi.org/10.1016/0375-9601\(91\)90116-P](http://doi.org/10.1016/0375-9601(91)90116-P)
- Valentini, A. (2010). De Broglie–Bohm Pilot-Wave Theory: Many Worlds in Denial? In: Saunders, S., Barrett, J., Kent, A., & Wallace, D. (Eds.). (2010). *Many Worlds?* Oxford University Press, 476-509.
- van Fraassen, B. C. (1980). *The Scientific Image*. Oxford University Press.
- Vickers, P. (2014). Theory flexibility and inconsistency in science. *Synthese*, 191(13), 2891–2906. <http://doi.org/10.1007/s11229-014-0464-8>
- Wallace, D. (2012). *The Emergent Multiverse*. Oxford University Press.
- Wallace, D. (2018). The Necessity of Gibbsian Statistical Mechanics. In submission.
<https://dornsife.usc.edu/assets/sites/1045/docs/gibbsboltzmann.pdf>
- Werndl, C., & Frigg, R. (2015a). Reconceptualising equilibrium in Boltzmannian statistical mechanics and characterising its existence. *Studies in History and Philosophy of Modern Physics*, 49, 19–31.
<http://doi.org/10.1016/j.shpsb.2014.12.002>

Werndl, C. & Frigg, R. (2015b). Rethinking Boltzmannian Equilibrium. *Philosophy of Science*, 82(5), 1224–1235. <http://doi.org/10.1086/683649>

Wittgenstein, L. (1953). *Philosophische Untersuchungen*. Blackwell.

Wu, Y.-S. (1984). General Theory for Quantum Statistics in Two Dimensions. *Physical Review Letters*, 52(24), 2103–2106. <http://doi.org/10.1103/PhysRevLett.52.2103>

Zurek, W. H. (2002). Decoherence and the transition from quantum to classical – revisited. *Los Alamos Science*, (27).