This electronic thesis or dissertation has been downloaded from the King's Research Portal at https://kclpure.kcl.ac.uk/portal/

Learning About Time from the Inside Out

Mason, Lucy

Awarding institution:
King's College London

The copyright of this thesis rests with the author and no quotation from it or information derived from it may be published without proper acknowledgement.

END USER LICENCE AGREEMENT

Unless another licence is stated on the immediately following page this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International licence. https://creativecommons.org/licenses/by-nc-nd/4.0/

You are free to copy, distribute and transmit the work

Under the following conditions:

- Attribution: You must attribute the work in the manner specified by the author (but not in any way that suggests that they endorse you or your use of the work).
- Non Commercial: You may not use this work for commercial purposes.
- No Derivative Works - You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you receive permission from the author. Your fair dealings and other rights are in no way affected by the above.

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 25. Dec. 2024

Learning about Time from the Inside Out

Lucy Mason

KING'S College LONDON

PhD Thesis
March 2024

Abstract

I argue that we should understand the flow of time using three ideas: information, emergence, and perspectives. Doing so produces an account that has a strong naturalistic basis but does not rely on fundamental physics. Instead it focuses on the theories that describe the realm of our experience; this allows us to justify our intuitions about time but removes the worry of subjectivity and avoids reducing flow to a psychological illusion.

I will develop a metaphysical account of time's flow that defines it as the localised becoming of events and argue that this should be understood via a perspectival ontology. From inside of time the world is presentist and from the outside it is eternalist. The principal focus will be on the importance and ineliminability of the presentist perspective; it is this perspective that allows us to understand how we have an open future and a fixed past, which is a crucial element of flow.

I will show that the internal perspective, along with the open future and fixed past, is an essential element of physics despite common assumptions that physics is perspective independent. Evidence for this will be found by looking at the role that information plays in physics and showing how this is an important factor in how we understand and interpret theories such as thermodynamics and quantum mechanics. I will analyse a number of these instances and provide new ways of thinking about how they use information.

This will support the major conclusion of this project: that the fixity of the past is due to localised records that allow us to make inferences about it. Likewise the future is open because there is no epistemic structure which fixes it, only probabilistic claims. Epistemic concerns like this are frequently dismissed as trivially subjective; but I will show that they are, in fact, an essential element of many of our physical theories and have important metaphysical consequences that cannot be ignored.

In particular, this model will provide a way to understand how flow connects to the well established view that time asymmetry emerges out of fundamental physics, and the wider emergentist project that shows we live in a levelled reality.

Acknowledgements

My deepest thanks go to Eleanor Knox, whose tireless support and guidance have been invaluable. Your philosophical insight has continuously challenged me and trained me into the philosopher I am now; without your input this thesis would be very different.

I also owe a great deal to Alexander Franklin, who provided many insightful comments and discussions. My philosophical views have been greatly shaped by yours, and your thoughts and advice have been deeply appreciated.

My examiners, Alastair Wilson and Oliver Pooley, provided incredible feedback and advice.

I would like to thank Karim Thébault who supervised my Master's thesis, which is where I first explored this topic, and James Ladyman who introduced me to philosophy of physics during my MSci at Bristol.

Particular thanks also go to David Papineau and Katie Robertson for many interesting discussions, and to David Wallace and John Norton, with whom I had valuable conversations during the term I spent at the University of Pittsburgh.

I am grateful to the London Arts and Humanities Project for the doctoral funding that supported this project and my visit to the University of Pittsburgh.

Finally this would not have been possible without endless support from my family and friends: my parents who have inspired my interest in philosophy, my sister who is always there to chat, and all the friends who provided encouragement.

Contents

1 Introduction

1.1 Time and its Flow

The flow of time is one of the most immediate and inescapable aspects of our experience. It is hard to imagine time without this feature. Yet providing a fundamental metaphysical account of it has been wrought with difficulty; even the metaphor of time as dynamic and moving immediately starts a troubling regress concerning what time is flowing with respect to (if not time itself). What is more, an investigation into time's flow must also account for the demands of physics, an arena that has frequently been seen as highly unfriendly to the claim that time flows. Some aspects of physics seem to deny that flow is a coherent notion at all; at best it seems physics may just have no use for it. It has become common to accept that the flow of time is just a psychological illusion in a static universe.

But even if we cannot find flow at a fundamental level of reality we need not abandon the project of finding a coherent account of it or give up on explaining how flow fits into a complete picture of reality. Nor need we relegate the question to psychology to look for explanations of an illusory phenomenon. If we are to successfully provide a metaphysical account of time and its flow then we should tread the line of attending to physics while also drawing lessons from how we experience time (which is our main evidence for the phenomenon of flow). The way to do this is to focus directly on the places where these two things come into close contact. There are three elements to this: first, the ways in which we learn about the world through observation and inference, and how we use the information we gain to predict the future and retrodict the past. This forms a basis for how we experience time. Second, the limits that our perspective - our position at the present moment, at a certain point in space, and with certain resources available to us - puts on our abilities to do this. And third, our placement at an emergent, macroscopic level of reality. So to find flow in physics we should look to emergent, perspectival elements of physics where measurements and information play an essential role, and examine how these physical theories are used to predict and retrodict. Doing so can give us a handle on the features of the world that our subjective experiences are derived from, and identifies places where the flow of time plays an important objective role in physics. It will reveal that, far from being superfluous, flow can help us understand how different parts of reality are related. Experience tells us where to look, but flow can be found in the world described by physics.

When we try to pin down our experience of flow more precisely we find a number of features. Certainly we have different attitudes towards the future (we hope, fear, plan and predict) than towards the past (on which we reflect, remember, mourn and celebrate). These attitudes are dictated by the belief that the past is fixed, unchanging, and certain while the future is open, full of possibilities, and subject to our control. This feeling is such an intrinsic and unavoidable part of how we experience the world that we rarely have any cause to doubt that this is in fact a metaphysical truth of the world. This belief, and the accompanying attitudes, can largely be attributed to the knowledge asymmetry (at the very least these things are closely connected). When we look around us we find detailed and abundant information about past events while the future is left to uncertain prediction. The connection between the epistemic asymmetry and asymmetries of causation, counterfactuals, deliberation and agency has been widely recognised.

It is difficult, however, to turn this into a more precise definition of what we mean by the flow of time. But a rough gloss of what we might take it to be is as follows: reality changes from moment to moment and flow is the dynamic element that transforms one moment to the next. More specifically it transforms an indeterminate future into a determinate past. The difference between past and future is essential to the transformative aspect of flow; it is not simply that there is reality at one moment and another reality the next. From the present we anticipate what is coming and watch one of the possibilities being transformed into actuality. These connections between moments are lawlike and specify how systems can evolve over time; this distinguishes time from space.

The project of this thesis, then, will be to show how we can connect our metaphysical claims about the flow of time - along with the open future and the fixed past - to the emergent epistemic structure in the world that allows us to make inferences about the past and future from a limited present perspective. This is the motivation for the title of this thesis - *learning about time from the inside out*: from our temporally embedded position we learn about the past and the future by using our scientific theories to predict and retrodict, and this is where we should look to find the flow of time. The aim of this is to fulfil our intuitions about time and draw on our immediate evidence of temporal flow (our experience of it) but do so in a way that removes subjectivity from the equation and locates flow in the way that measurement and information are used in physics. The result is flow that appears primarily at a macroscopic level of reality; but appears as an objective and necessary element of reality. This account provides a middle path between abandoning flow to the realm of psychology or insisting that flow is found in our most fundamental theories. This is an option that has received little attention; there are no accounts of time's flow as an objective yet emergent phenomenon despite the well-established parallel project on how the asymmetry of time emerges from a fundamentally symmetric world.

For a metaphysical account of flow I will use the concept of *becoming*. Originating with Broad (1923), becoming is a metaphysical notion that is meant to capture the dynamic property of time while avoiding the more troubling implications that the word 'flow' brings with it; it is the coming into existence, occurrence, or happening of one moment after another. I will particularly look at *localised temporal becoming*, which is a more recent development of the view that extracts Broad's notion from the growing block context in which he introduced it and pairs it instead with a primarily eternalist ontology (although as I will argue this rough gloss is potentially misleading).[1] Using this as a basic notion of flow I will explore what becoming really amounts to and how it connects to the metaphysical claim that the future is open and indeterminate. I will argue that the past is fixed because there are records of it all around us, while the future is metaphysically indeterminate because it is only possible to make probabilistic claims about it. Flow comes into the picture when from one moment to the next more information becomes available and settles questions about the possible future.

---

[1] The updated view has come from a number of authors including Arthur (2019), Savitt (2002; 2006), Ismael (2016a), Dieks (2006), Dorato (2006).

The three elements that go into this project - information, perspectives, and emergence - each require justification. Of the first, information (which I take broadly to encompass general epistemic features as well as more technical definitions of information), I have already said that the epistemic asymmetry is closely connected to the way that we experience the flow of time and the beliefs we have about the past and the future. Epistemic claims like this are frequently dismissed when we are looking for metaphysical answers because they are taken to be trivially subjective; they tell us about our own abilities or ignorance and not the external world. Any accounts that do recognise a connection between information and time's flow relegate this to psychology. A common method for understanding how an illusion of flow arises is to model human agents as information gathering and utilising systems (IGUSs). The idea being that the way a physical system can use memory to deliberate and act creates the illusion of asymmetric flow. This has a basis in the physics through the limits on IGUSs that come from it being a physical system; but flow is a product of how the different parts of the system work together to create a continuous subjective experience.[2] Despite the physical basis for phenomenal experience the effects themselves are psychological and not metaphysical. However, properly understanding how epistemic claims are made through physics will show that the metaphysical implications of information should not be so easily dismissed. The process of learning about the past and the future from the present is such a commonplace part of physics that it is often overlooked. Likewise the structures in the world that can be used by us either to gain knowledge or process it are assumed to be of operational use, but their role in physics beyond that is neglected. But by acknowledging the role that information plays in both developing and interpreting theories we can see how flow is present in physics and it can help us understand the way our theories represent reality.

I will now look at emergence and perspectives in turn and show how both approaches connect to and support the use of information in interpreting physics, and how each may offer novel insights into the nature of time. The use of these

---

[2] See Callender (2017) for example; Callender as calls time the *great informer* due to how dynamical laws take us from past to future. Less explicitly, similar ideas are found in Ismael (2023). The use of IGUSs has also picked up traction in quantum mechanics where it can be used to model observers' experiences of measurement - although not in connection with time's flow (Gell-Mann & Hartle 1994).

concepts has become more common and they have been increasingly well developed; yet they remain controversial in many ways.

1.2 Naturalised Metaphysics and Emergence

Flow is first and foremost given to us as an aspect of our experience. If we want to avoid subjectivity then paying careful attention to the physical processes involved will objectify the intuitive insights we might gain from our everyday experiences. The highly influential program of naturalised metaphysics has discredited an appeal to intuition. Intuition is too subjective and is moulded by the limited portion of reality that we have access to. Our personal experience cannot capture the lessons that science teaches us about a world that extends far beyond us. We should abandon intuition in favour of science. As Ladyman and Ross put it:

> "Since scientific institutions are the instruments by which we investigate objective reality, their outputs should motivate all claims about this reality, including metaphysical ones." (Ladyman & Ross 2007, p. 30)

Sticking closely to science and showing how all metaphysical claims are motivated by at least one piece of physics (where physics is prioritised due to having the most universal scope of all the sciences) brings scientific rigour to metaphysics.

However, Ladyman and Ross also strongly advocate that we recognise that not all of science is fundamental physics. We use a vast array of higher level theories both within physics and in the other sciences that offer novel insights into the nature of the world. These theories have just as much methodological rigour and frequently identify structure in the world that is missed by lower level theories. This is counter to the strong reductionist view, which I take to be the view that all of science reduces to fundamental physics and only fundamental physics can tell us about the ontology and foundational nature of reality; higher level theories are useful but merely pragmatic abstractions or approximations to physics. Naturalised metaphysics has revived the program of emergence and provided compelling support for taking higher level theories seriously when it comes to drawing metaphysical conclusions. The resultant emergentist metaphysics argues that, while science is unified and it remains important to examine the intertheoretic relations between levels (in this sense emergence is compatible with a weaker and non-eliminative form of reduction), we should accept macroscopic objects into our ontologies and recognise the explanatory value of higher level theories.[3]

Emergent theories are particularly useful for explaining our experience. Even if naturalised metaphysics has discredited appealing to experience-based intuition when doing metaphysics, it remains important to give an account of the level of reality that we have direct access to. The world of our experience is a legitimate level of reality; we do not experience a world of fundamental physics but an everyday macroscopic reality governed by theories such as classical mechanics and thermodynamics.[4] When it comes to explaining our experience of time, then, it is these theories that we should look to. The approach described here satisfies naturalistic metaphysics by closely relating all its metaphysical claims to relevant aspects of physics, even if it uses experience as a guide to which aspects of reality to focus on.

So far the most thorough application of physics to questions about time is the debate on time asymmetry.[5] Alongside projects attempting to find the direction of time in fundamental physics, there is a highly successful program where asymmetry emerges at higher levels (most commonly in thermodynamics and connected areas such as quantum decoherence). This has been developed in great detail and connected with other asymmetries such as those of causation, counterfactuals, and agency. This program is both reductionist in that it aims to unify all the temporal arrows by reducing them to thermodynamics, but also emergentist in the way it describes how the asymmetry of thermodynamics emerges out of a fundamentally symmetric theory. It shows how a distinct and substantive phenomenon - time asymmetry - appears at higher levels. Asymmetry is ineliminable and highly

---

[3] The emergentist program did not begin with Ladyman and Ross but their work has been influential in rejuvenating it. The emergentist program has received increasing attention and has now been substantially developed. Emergence has also come to mean something different in the philosophy of physics literature than it does in, for example, philosophy of mind. It is a weaker concept much more closely related to reduction.

[4] Even then there may be a disconnect between formal scientific theories and a folk theories of the physical world.

[5] Of course much of this predates the formal campaign for naturalised metaphysics by Ladyman and Ross, but they make the point that even where physics *has* been applied to metaphysics it has not gone far enough.

important for our understanding of science, but not universal or exceptionless.

The ontology and flow of time are comparatively much less discussed; the only major and widely agreed upon application of physics to this has been the challenge that special relativity makes for defining simultaneity and the present moment. More recently quantum mechanics has been brought to bear on this issue with discussions of quantum indeterminacy and the possibility of quantum becoming (e.g. see Callender 2017). This work takes a fundamentalist approach to flow, for example attempting to model it using theories of quantum gravity such as causal set theory (Grandjean 2022). Little work has been done to connect the emergentist project explaining the arrow of time to debates about flow.[6] But, while passage *could* be described as symmetric, our most intuitive understanding of flow implies flow in one direction or another. Putting these two debates together and seeing both direction and flow as emerging side by side will provide a unified picture of time.

What is more, the way that different theories are used to make inferences and predictions is closely related to the emergentist project. Different levels of reality are described by different theories and have different resources available for these tasks. For example the introduction at higher levels of features like probability and statistical methods significantly change what our predictions look like. This means that using information to understand physics goes hand in hand with considering emergence. An emergentist picture of time's flow will take into account the full spectrum of physical theories at different levels and how they deal with time in different ways.

The importance of the epistemology of science for metaphysics has not gone entirely unrecognised. When presenting naturalised metaphysics Ladyman and Ross emphasise not only paying attention to physics but also the methodology of scientific institutions in extending knowledge claims. This is even more clear in the perspectival realism project (Massimi (2022) gives its most thorough and up to

---

[6] Maudlin (2007) does connect the passage of time to thermodynamics but remains committed to a fundamentally reductionist program. The connection only follows because flow is committed to a direction. Arthur (2019) gives similar considerations where higher level asymmetry must follow from fundamentally asymmetric becoming.

date presentation). Here realism about science is built around the epistemological project scientists engage in; our individual experiences are built into collective knowledge through the scientific method of observation. Both these accounts are primarily putting forward distinct ways of understanding scientific realism; I will not engage with their specific proposals on this subject. However, neither really looks at how information is used within physics itself. (Ladyman and Ross frequently mention information theoretic ideas but more as a methodological process than as a physical aspect of the world to study.) What they really offer is a guide for how to do metaphysics more generally: what is discoverable and through what means should be the starting point for metaphysics. But, instead of looking at the historical and social aspects of science as a collective project, I will take this guide much more literally; metaphysics can be derived from how physics uses and provides information. Formal theories such as information theory as well as general considerations about measurement and predictability make this precise and objective.

1.3 Perspectival Physics and Perspectival Metaphysics

The latter part of the title of this thesis - from the inside out - deserves more unpacking. Central to the arguments I will make are the epistemic methods we apply to learn about the past and the future *from our own limited perspective.* We are trapped inside of time and can only ever experience the present moment. It is due to this perspective that we are forced to make complex epistemic inferences about other times in our attempt to understand the world. In contrast to this, physics is often taken to be objective and independent of any observer. This is seen as the goal of science - to get rid of the experimenter and describe the world ‘as it really is’. As a result it is often assumed that physics supports the eternalist view of time (where all times are equal) and is an obstacle for presentism (where only the present exists). Presentism requires you to index facts in relation to the present, and singling out some particular privileged moment through physics is notoriously difficult. Part of the view of temporal becoming that I will lay out as an account of time’s flow will be the idea that presentism and eternalism are two different perspectives on time; presentism is the perspective from inside of time and eternalism the perspective from outside of it.[7] Showing that perspectivalism does not equate to subjectivity and that internal perspectives have explanatory power that the external perspective does not will be a large part of this thesis.

Why should we take perspectives seriously, and what kind of perspectives are we using? The use of perspectives has been gaining momentum within philosophy of science for some time. We now have well-defended positions such as causal perspectivalism (e.g. Price 2007) and many different perspectival interpretations of quantum mechanics (ranging from QBism to Relational Quantum Mechanics - the latter will be considered in chapter 6) - as well as arguments for the fundamentality of open systems even where they have observer or perspective dependent boundaries (Cuffaro & Hartmann 2021)[8]). There are non-perspectival alternatives to all these positions, and they remain contentious, but they show that the premise is coherent at the very least. Price summarises this nicely by saying that perspectival reasoning asks the question "What kind of reality would look like this, from the particular standpoint we humans happen to occupy" (Price 2007, p. 290-291). The idea being that our account of external reality should explain why we experience perspectival phenomena the way that we do.

All these positions come with careful justifications about how perspectivism can deliver objectivity, although they range in how anthropocentric the definition of perspective is. Some examples of perspectivism are explicitly subjective (for example the agents used in QBism). I want, however, to focus on objective and physically defined perspectives. An initial example to start with, in which a debate about perspectives (of a sort) is already playing out, involves spacetimes and the importance of differential geometric versus coordinate-based approaches. This debate is certainly not taken to involve anything subjective. We can understand coordinate systems as perspectives in the following way: at its most simple a coordinate systems indexes claims about spatial and temporal relations to a reference position and form of motion. Most often we use physical objects to mark different

---

[7] This idea comes from Savitt (2006). Chapter 2 will explore this in detail.

[8] Cuffaro and Hartmann write in relation to interpreting open systems and comparing them to Copenhagen interpretations that "the ideal of an observer-independent reality is not methodologically necessary for science and that modern physics (especially, but not only, quantum theory) has taught us...that there is a limit to the usefulness of pursuing this ideal." (2021, p. 23)

coordinate frames; in pedagogical examples of special relativity we often talk about a rocket's coordinate frame versus the earth's frame. An object's own rest frame often has special significance, such as defining the concepts of proper length and time. The object is physically embedded in space and time and its position - i.e. its perspective - is worked into the physics. Of course, in more sophisticated examples (such as generalised coordinates) the idea of coordinate systems diverges from this basic understanding of an embedded perspective. However, all maintain a sense of defining the coordinates with reference to something specific.

The coordinate-based approach indexes all the claims of the theory to a reference frame and focus on equivalence classes of reference frames; in contrast the differential geometric view focuses on invariant mathematical structure. But Wallace (2019a) argues that both formulations peacefully co-exist in the physics literature and the eschewal of the coordinate-based approaches in foundational philosophical discussions is a mistake. There are a number of advantages that the coordinate-based approach has in certain circumstances. First, it is common practice in physics and simplifies many problems considerably:[9]

> "While any such theory can usually be cast into a coordinate-free form, doing so can often be complicated and can render obscure pieces of physical reasoning that are fairly transparent from the coordinate-based perspective." (Wallace 2019a, p. 132-133)

Second, coordinate-based approaches help understand absolute structure in spacetime theories. For example see Weatherall's (2016) discussion of Maxwellian spacetimes compared to Newton-Cartan theory and how the identification of empirically privileged reference frames (i.e. ones matching trajectories that material bodies actually follow) indicates absolute structure (see also Wallace 2020).[10]

---

[9] Wallace also points out the pedagogical advantages of the coordinate-based approach. "True mastery also requires easy familiarity with coordinate-based approaches." (Wallace 2019 p. 135)

[10] Another example that Wallace (2019a) uses of where the coordinate-based approach may have more explanatory value is Brown's (2005) argument that the symmetries of the dynamical laws (which capture coordinate transformations) are more explanatory than the geometry of spacetime. Brown says: "That *is* the prima facie mystery of inertia in pre-GR theories: how do all the free particles in the world know how to behave in a mutually coordinated way such that their motion appears extremely simple from the point of view of a family of privileged frames? To appeal, however, to the action of a background space-time connection in which the particles are immersed...is arguably to enhance the mystery, not to remove it." (Brown 2005, p

For the subject of spacetime, then, the perspectival coordinate-based approach has clear potential value despite being inherently indexical, and there is no confusion that relativity is subjective or observer dependent. Observers are used as a shorthand for references frames. The coordinate frames, and even more importantly the transformations between them, capture important structure of the theory. We certainly should not give up the differential-geometric approach but neither should we be scared off by the perspectivalism inherent in reference frames. Similarly, I will apply this to time to argue that even if the eternalist picture is coherent we should not dismiss presentist perspectives. However, the reasons Wallace gives for taking coordinate-based approaches seriously largely focus on how they are *just as good* as the differential-geometric approach and sometimes pragmatically more useful. But I will go further than claiming just that we should take perspectives seriously and also say that they are ineliminable. There are aspects - such as flow - which a non-perspectival view of time cannot capture at all.

Coordinates frames only capture one element of perspectives (a physically embedded position that is used as a reference). The notion of perspectives I will use goes beyond this, and this is where perspectives connect with information: a perspective is defined by the information available from it. Remaining in the spacetime literature, a minimal notion of this is captured by the lightcone structure, which places a constraint on what events someone in the present could know about, and is often discussed as an important limit on causation.[11] However, we can go beyond this to identify many other ways in which the information available in a perspective is constrained. What I want to retain from the example of coordinate systems is their objectivity. An observer might be used as the reference to define a coordi-

---

142) Instead he advocates explaining relativistic phenomena in terms of the dynamical laws the particles individually follow; spacetime is a geometrical abstraction of the symmetries between coordinate systems found in the laws.

[11] Another area in the literature on spacetime can also be used to illustrate the importance of observability and practical methodology for understanding the structure of the world. Theories of spacetime are often built around dynamical symmetries and one of the best ways to understand symmetry relations is to look at what structure in the theory is (in principle) observable and what is unobservable (Wallace 2019b; Greaves & Wallace 2014). Understanding what a measuring device is and how features of a theory are rendered observable is a guide to identifying and interpreting dynamical symmetries. (However, the notion of measuring device used is very abstract and does not relate to practical measurements, the argument rests on how any device is itself bound by the symmetry.)

nate system but the model of the world that is picked out through it does not depend on the observer, the observer is just a useful placeholder to understand the perspective.

What information is available from a perspective depends not just on our embedded position but also on the mechanisms of measurement and observation.[12] The measurement process has become undeniably relevant to the foundations of physics when it comes to quantum mechanics. The measurement process, including what information can be extracted from quantum states and how the physical apparatus of measurement should be understood, is an integral part of all the different interpretations of quantum mechanics. More and more we have come to recognise that observability separates out which elements of a theory we can make well motivated metaphysical claims about. Going beyond the empiricist position, this considers not just observational data but also the process of measurement itself. This sort of post-positivist approach looks at how observations are made, in what context, and how they fit into the overall theory; a formal theory of measurement is used as a tool to help us understand the theory itself and what inferences can be drawn from it.

Defining perspectives in terms of information and observations is a useful way to identify the use of perspectives in physics. How systems gain, lose, transfer, and transform information can tell us a lot about the structure of the world. Despite the tendency to dismiss this as non foundational, operational limits on information and observability are frequently a sign of something deeper. Either they highlight important structure in a theory or they can indicate the importance of *non*-fundamental theories. Thermodynamics, for example, has long been accused of being subjective or non-fundamental due to its information theoretic underpinnings; there is extensive debate about whether the probability distributions that it uses are a result of our subjective inability to precisely measure the microstate

---

[12] The use of perspectives here is most similar the Price's causal perspectivalism, in which a human perspective is defined by what information we have available and what ways we can act on the world around us. However, I want to place far more emphasis on the *physicality* of this, and not our role as agents in this perspective. It is the physical structure of information and measurement that define the perspective under my usage, and not the agent which holds the information. Chapter 3, and the discussion of Ismael's (2023) view of the open future, explore this in depth.

or if they represent something more objective.[13] But the expansive literature on emergence and reduction - as discussed in the previous section - has challenged the dismissal of higher level theories and shown that they have novel explanatory value that is not delivered by the reducing theory. Emergent theories are important examples of a sort of perspectivism (although it may not be common to label them as perspectival). Some theories describe the world from a macroscopic perspective in which microscopic information is either difficult (or impossible) to obtain or flatly irrelevant; this comes with a certain set of theoretical tools and resources which are appropriate for that perspective. Other theories give the microscopic perspectives. And well-defined relations of emergence and reduction lay out how we can transform between these perspectives just as Lorentz transformations do for coordinate systems. Recognising the restricted domain of a theory outlines a perspective in which that theory is valid; part of what this thesis will explore is how exactly emergent theories relate to the limits on information set by perspectives. This will focus on how emergent theories are particularly valuable for making inferences when we occupy a perspective with incomplete and localised access to the world.

All these uses of perspectives have in common that they index the use of the theory to some position, from which a theory is deployed to make inferences about what will happen in the world based on the information available in that perspective. The world may appear different from another perspective, and a different subset of information will be available, but this does not negate the need to understand each perspective in its own right. How the perspectives relate to each other highlights important structure in physics that unifies the perspectival phenomena; yet the novel contributions of each perspective are distinct.

This, I hope, goes some way to justifying the use of perspectives and illustrates how exactly they can give us novel and foundational (even if not always fundamental) ways of understanding the world. The assumption that physics is perspective free is mistaken and the erroneous focus on the 'view from nowhere' obscures important facts about the world. We see this mistake copied when it comes to time in the assumption that the determinate eternalist reality explains all we need to know. The arguments presented in this thesis will pick up on many of the strands of

---

[13] See Robertson & Prunkl 2023 for a discussion of this.

argumentation used in the examples just given. Imagining time as if we are looking down on it from above and it is entirely separate from us leads us to overlook many of its features that may only be apparent from the perspective embedded in time that we actually occupy. These features are indexed to this perspective but are no less objective for it. Neither are they subjective or psychological; they are physical aspects of the world itself independent of us. How we use these physical features to learn about other times from our limited perspective, and our differing epistemic access to the past and the future, can teach us about the nature of time. Flow is integral to the way that the world presents itself to *us*, so to locate it in physics we should look for the physics that governs the world from our perspective. And in doing so we find that epistemic structure has unique explanatory value.

1.4 Plan for the Thesis

I have motivated the methodology and overarching questions of this thesis. But how exactly will I argue for it? My aim is to show how localised temporal becoming, understood in terms of a perspectival ontology, provides an account of flow and the indeterminateness of the future. I will show that this should be directly linked to emergent epistemic structures in physics such as future probabilities and physical records of the past. The result is that, although becoming occurs at all levels of reality, it only forms a robust account of flow at this emergent level. First I will lay out the metaphysical half of the project (in chapters 2 and 3) and then dive deeper into the physical justification of its more contentious claims by looking at several areas of physics that use information and records in relevant ways.

My aim is not to definitively identify all the places we can find relevant structure or to tie this model of time to specific theories. Instead I will consider a range of theories and interpretations that all use information in interesting ways and support the various aspects of this project. Doing this will make it evident that the use of information (or more generally epistemic structure) in physics is extremely widespread; thereby demonstrating that this approach to understanding time is widely applicable and can be pursued regardless of exactly what theories we commit to. The use of information in multiple areas, each found to support the claim that the future in unknown and indeterminate while the past is known and fixed, is evidence that this view is a robust and fruitful interpretation of physics. It will also repeatedly illustrate the close links between information, emergence and perspectives. As a result I will frequently draw on theories that are not necessarily compatible with each other; for example I will extensively use ideas from special relativity about the lightcone structure alongside quantum mechanics without delving into relativistic quantum mechanics in any depth (although none of the aspects of quantum mechanics I use rely on having a preferred reference frame). I will also use both quantum and thermodynamical probabilities without assuming any particular connection between them. And finally I will discuss the range of interpretations of quantum mechanics without committing to any particular view. The value of looking at such a broad range of theories is that it justifies the methodology I have set out here and allows for a model of time that depends on generic, recurrent, features of the world rather than specific theories that may be subject to change.

I will now give brief summaries of each chapter to lay out how the argument will go:

Chapter 2 - Becoming the Block

The purpose of this chapter is twofold. It will serve as an introduction to localised temporal becoming and lay out how it fits into the literature on the ontology and flow of time. This will involve explicating key features of the view that are an important foundation for later chapters, such as its commitment to localised flow and a non-standard present. But the more incisive purpose of this chapter will be to argue that localised temporal becoming is best understood in terms of perspectives. Presentism is the perspective inside of time in which only the present exists (in a tensed sense of existence), and eternalism is the external perspective in which all moments exist (in a tenseless sense). This has been proposed by Savitt (2006) but is often not explicitly recognised. Understanding the view in these terms - and showing that the internal tensed perspectives are ineliminable - helps defend it against objections that localised temporal becoming is trivial and uninteresting. I will also argue that, based on this understanding of the ontology, localised temporal becoming is compatible with the claim that the future is open from the internal, tensed, perspective. This is necessary for its account of flow and further strengthens the view.

Chapter 3 - The Open Future and the Recorded Past

Having laid out the ontology the task remains to show that the internal perspective - and the open future - are indispensable if we are to provide a complete description of objective reality. I will show that what it means for the future to be open is for claims about the future to lack determinate truth values because the present state of the world combined with the laws is not enough to predict with certainty. I will then turn to what constitutes a failure of predictability. Current views either look for strong indeterminacy in quantum mechanics or accept the indeterminacy to be merely epistemic and focused on deliberation and agency (e.g. Mariani and Torrengo 2021a; Ismael 2016a, 2023). I will argue for a middle ground where we recognise the physical basis of information (for example, the lightcone structure of spacetime, probabilities in higher level theories, and records of the past) as an objective feature of the world that leaves the future metaphysically indeterminate in a perspectival way. Finally, I will look at the counterpart of the open future - the fixed past - and argue that fixity must be dependent on records of the past. Understanding both openness and fixity is essential if we are to understand the transformative aspect of time's flow.

Chapter 4 - A Non-Idealised Account of Records

The remaining three chapters will focus on examining the physical basis that I rely on in previous chapters and on demonstrating the role of information - particularly records - in physics. I will focus on records as there is relatively little analysis of records compared to the well established literature on probabilities and prediction. The most well known attempt, from Albert (2000; 2015), to explain what they are is highly idealised and abstracted from their practical use. This does not match how records appear in our physical theories where taking account of the details Albert idealises away is essential to our theoretical development. A notable example of this, discussed in this chapter, is the use of Szilard's thought experiment, the one molecule gas memory device, which has been extensively used in attempts to prove Landauer's principle. I will present a new account of what records are, building on Albert's basic construction but explicitly considering how they work in a complex and interconnected world. This new definition, which focuses on the shielding off of noise, allows us to better understand the use of records in quantum computing and the thermodynamics of information processing as well as providing a general account of how naturally formed records are informative. It also establishes that records are primarily a feature of the macroscopic world.

Chapter 5 - The Role of Records in Quantum Decoherence

Quantum Darwinism explains the emergence of the classical world from the quantum by modelling decoherence in an information theoretic framework. It identifies the key criterion of classicality as the possibility for objective knowledge of the state shared by many observers. This differs from the more commonly used standard for emergence in the philosophy of physics literature that relies on the instantiation of classical dynamics. Quantum darwinism argues that intersubjective agreement is achieved when an interaction between the system and the environment produces redundant records of the state in the environment. The goal of this chapter is to examine the use of records in quantum darwinism; I will use the account of records developed in the previous chapter to explain why they are so critical in quantum darwinism, and how understanding them allows us to relate the emergence described by quantum darwinism to general philosophical accounts of the emergence that focus on dynamical laws. Doing so reveals how records are part of (and indicators for) the level dependent structure of reality and provides a template for how we can understand emergence in terms of information.

Chapter 6 - The Propagation of Definite Values in Relational Quantum Mechanics

This chapter will consider relational quantum mechanics as a case study for how the metaphysical picture of time I have developed can be applied to a specific theory. Although I do not want to tie this account of time solely to this interpretation, relational quantum mechanics - as an interpretation of quantum mechanics formulated entirely in terms of the information available to different perspectives - is a useful theory in which to demonstrate how we can understand the world in terms of the model I have proposed. I will also show how the connection between temporal becoming and records can suggest solutions to some current challenges facing relational quantum mechanics, namely the problems around coordinating information over time and across perspectives.

Finally, in the conclusion I will draw together all the elements and show how the exploration of the physics in the later chapters ties into the metaphysics of the early chapters. This will particularly spell out the role that emergence plays in the final picture and how this affects our understanding of the relation between presentism and eternalism at different levels. The resulting picture claims that becoming is a foundational part of reality but what makes becoming a substantive account of flow is the emergent epistemic asymmetry found in physics. Moments are connected by dynamical laws - as well as other techniques of prediction and retrodiciton - that tell us how what exists in the present moment will change. When, at the higher level, these relations display asymmetry then they give the sense that time is moving forwards and transforming an indeterminate future into a fixed past. Different levels of reality use different theories and have different access to information; this way of understanding time is the best way to think about time's flow in a levelled reality. The emergence literature has pushed us to accept that there are higher level ontologies that do not exactly map onto fundamental physics. Continuing to treat time as flat and defined only by a single level is non-functional.

2 Becoming the Block

2.1 Introduction

This chapter will examine a family of similar views - that I will call localised temporal becoming - which identify time's flow as the coming into existence of events along a worldline. The phenomenon (flow) that it attempts to capture is elusive and has been defined in many different ways. What can be identified as a shared intuition is that flow is a dynamic process that relates one moment in time to the next. There are two parts of this that an account of flow needs to grasp: a) that reality changes from moment to moment and b) what the relationship between moments is that goes beyond a simple ordering relation. The first part is easier to describe. It is the sense that the world changes over time, either it is constituted by different facts or a different state of affairs obtain; there are many possible ways to express this. The second part is harder to unpack. There is not just change but also well defined relations between the world at one moment and the world at the next. Instead of just an ordering of events it seems that one moment is *transformed* into the next. The categories of past, present, and future are vital to this. Time passes because the future is transformed into the present and then into the past.[1] This distinguishes time from space; there are changes over space - the valley rises to a mountain - but not in this transformative sense. I will show how localised temporal becoming can capture both these elements.

The exploration of localised temporal becoming so far has not clearly articulated its ontological basis or shown how its understanding of existence contributes to its account of flow. This has led to confusion and a number of criticisms. I will argue that the strongest interpretation of this view takes it to be based on the claim that the difference between presentism and eternalism is perspectival. Presentism is the perspective from inside of time and eternalism from outside of it. The eternalist aspect of this account has been overemphasised and this has led to a neglect of the important role that the presentist perspective plays. The latter should be

---

[1] This also comes with a sense of direction that I will largely put aside for now and return to in chapter 3.

considered to have primacy when it comes to explaining flow. It explains how a theory with an eternalist block can maintain that there is a genuinely open future from the presentist perspective, and that becoming is the process of transforming future possibilities into actuality. This is essential to understanding how becoming captures the transformative aspect of flow that is central to our experience of time. Our experience is unavoidably located *inside* of time and so it should be expected that the embedded perspective is where we find this phenomenon.

I will first lay out what temporal becoming is and how it is being used in this family of views. Then I will consider a range of preliminary arguments from both physics and philosophy, many of which are standard in the literature, to see how these motivate this position (and to set up features of the view that are relevant for future chapters of this thesis). This will also illustrate why many of the other well-known models of time fail. In particular this view is able to overcome difficulties that discredit standard presentism, the moving spotlight, and the growing block. With this firm basis I will then turn to looking at how we should treat existence with respect to time and why understanding a perspectival relationship between presentism and eternalism helps to defend this account of passage. Finally I will show that the open future, which is often cited as the hallmark of temporal flow, follows from this model.

The purpose of this chapter is to argue that localised temporal becoming provides a coherent and meaningful account of both the ontology of time and its flow and to emphasise the importance of perspectives in making sense of this. This will largely focus on a fairly minimal notion of perspective at this point, using mainly the idea of an embedded position. The connection between information and perspectives will start to be explored when considering the open future, but the full connection will be made in the next chapter. Understanding time as a series of tensed events woven together by dynamic temporal becoming gives a strong basis to further explore the role that the open future plays both in our experience of time as agents embedded within it and in the physics that we take to describe the world. This exploration will continue into the next chapter where I will take a closer look at what the open future amounts to, how it appears in physics, and how it connects to ideas about information and emergence.

2.2 Temporal Becoming

Temporal Becoming originated with Broad (1923) in the context of the growing block theory of time to capture the idea that each moment is continuously superseded by the next as new moments are added to reality. It is also commonly referred to as a moment (or event) *occurring* or *happening.* This has been largely interpreted to describe time's flow as it is a good candidate for capturing the sense that reality changes from moment to moment.

Broad's becoming was strongly linked to existence. The growing block describes how a slice of reality is added at each moment, forming a block of time where the present and past exist but the future does not. The present is defined simply as the moment that has no moments after it. When the next slice of reality comes into being it shifts the present onto the newest (and latest) moment. Broad's later work moves away from the growing block theory somewhat into a general relational view of time (see Thomas 2019 for a historical account) but he retains the concept of becoming, although it becomes less clear how it links to existence. Since then temporal becoming has been picked up widely in the literature on time and used in numerous different ways. In particular there has been a new conception of it as a way to add dynamic flow to an eternalist block universe conception of time. A number of different authors have made similar proposals (Arthur 2019; Dieks 2006; Dorato 2006; Rovelli 2019; Deng 2013a; Deng 2013b; Ismael 2016a; Leininger 2021; Oaklander 2015; Savitt 2002; Mozersky 2015 - among others). The overarching idea of these accounts is that there is no privileged present and that all of time exists in a block universe. And yet dynamic becoming occurs at each moment, capturing the transition from one moment to the next. This view has largely kept the name temporal becoming but has also been called tenseless passage (Deng 2023) and temporal B-coming (Leininger 2021). Another name for it is localised temporal becoming (Arthur 2019), this comes specifically from versions that consider the tension with special relativity and localise becoming to a worldline. (Section 2.3.1. will consider localisation in detail.) Localised temporal becoming is the name I will use throughout this thesis as it avoids the troublesome focus on the block universe aspect of this view (more on this to follow). I will note where temporal becoming is taken to mean the original sense in the context of the growing block.

The response to this parcel of views has not been universally favourable. Earman (2008, p. 159) calls this sense of passage "thin and yawn-inducing" and Pooley (2013) similarly dismisses it as not saying anything that is not already obvious and widely agreed upon. Saying that events within the block *occur* in a temporal sequence is trivial and gives no meaningful notion of time's flow.[2] This is especially true when becoming is taken out of the growing block model and placed in a context where there is no difference between past and future; the transformative aspect of flow seems to be lost. This notion of becoming is too insubstantial to add anything to a block universe.

As I will explore in this chapter, these criticisms largely stem from divorcing the concept of becoming from questions about existence. The original concept in the growing block theory was given its metaphysical weight by the claim that a moment came into being and added to the set of things that exist. This was baked into the ontology of the growing block account and its denial of the existence of the future. Becoming made actual this previously non-existent future (fulfilling both the change in reality and transformative requirements of flow). However, with the shift to applying becoming to a block universe, becoming was used much more liberally. The connection with existence becomes somewhat muddied and the sense of a previously non-existent future being materialised no longer seems to apply. Becoming is reduced to events *occurring* with no real description of what that is supposed to mean. There is diversity in the family of views of localised temporal becoming about how closely they stick to the block universe conception of existence. Pooley takes the position to be committed to the fact that the tenseless block universe, which catalogues all the events and the temporal relations between them, is all there is to say about reality (i.e. everything exists in the block universe and this is the entirety of reality). Becoming is then an addition to this that does not latch onto anything meaningful about the nature of time; the future has exactly the same status as the past in the block universe. Certainly some authors in this group deny that becoming has anything at all to do with ontology. Leininger (2021) is an example of this, although she connects becoming with causal powers

---

[2] The particular authors they direct this too are Dieks (2006), Savitt (2002), and Dorato (2006) which they take as examples of the general view.

instead to try to give it the necessary metaphysical weight. But this is a minority view, most do explicitly retain some connection with existence. They reformulate becoming in terms of a new understanding of the presentism/eternalism debate and conceive of existence in a fundamentally different way (Rovelli 2019; Savitt 2006; Dieks 2006; Dorato 2006). For these authors Pooley's strict characterisation of their views as committed to an essentially B-theoretic worldview is mistaken. Section 2.5 will explore this further and it is ultimately the conclusion that I will argue for. Finally, many versions of this view are presented in terms of tensed and tenseless facts and the A-B debate stemming from McTaggart (Arthur 2019; Ismael 2016a; Deng 2013b). While closely analogous to presentism and eternalism - and questions about existence - these distinctions can come apart. But I will largely take these views to be in agreement with the camp that adopts a new understanding of the relationship between presentism and eternalism, although these authors do not make this explicit.

There has been no systematic examination of these views and what they agree on. I will not attempt to do one here and instead will proceed by drawing elements from many of the views to present the strongest possible account. Doing so will show that these views should not be understood in terms of the basic block universe and that their nuanced analysis of existence can accommodate a non-existent future coming into being. This reinstates temporal becoming as a substantial metaphysical process and responds to the worry that it is a trivial addition. I proceed with the acknowledgement that individual authors may disagree with this reading.

Despite the differences between these views when it comes to existence and the general unclarity on this issue, there is a clear agreement on the general features of this account of time's flow. There is a consensus on a localised version of becoming being necessary (the becoming of spatiotemporally localised events rather than universal moments). The idea of a single privileged present moment is abandoned and all times are treated as equal. And the block universe is taken to be an accurate, although not necessarily complete, representation of time. Before looking more closely at the connection to existence, and what this means for defining becoming in terms of the future coming into being, it is worth justifying these features independently and making a comparison to see how this account of temporal becoming fares against other prominent accounts of time's flow.

2.3 Preliminary Arguments

This section will present three standard arguments, and one less standard one, which commonly crop up when considering time's flow. These have been much debated and will be familiar to most readers but I will take the time to lay them out here in some detail. Doing so will provide motivation for the common features of accounts of tenseless passage mentioned in the previous section. These features will be relied on in later chapters of this thesis so it is important to lay them out thoroughly here. This discussion will also provide a comparison to the common models of time's flow that are defeated by these problems. (The other models I will mention are standard presentism, the moving spotlight, and the growing block models of time.[3]) Laying this out explicitly will form a clear foundation on which to explore the more nuanced arguments in favour of localised temporal becoming. It will also set up the arguments that will, in section 2.4, prove to be fatal to the growing block view. The failure of the growing block clears the stage for a more careful consideration of the connection between becoming and existence.

2.3.1 Universal versus Local Flow

We commonly think of moments in time as being defined universally by the relations of simultaneity between events. From this follows the idea that becoming proceeds from one universal moment to the next. However, special relativity has shown that simultaneity is relative to a reference frame and there does not exist a unique foliation of spacetime. B-series ordering relations of earlier and later can be partially retained under special relativity by limiting the order to inside the light-cones of events but a universal procession of time cannot be established. Moving to general relativity does not help with this either. Earman (2008) discusses the

---

[3] This is not an exhaustive list. Another account is the open futures model that will be discussed later in this chapter. There is also flow fragmentalism. This has many similarities to localised temporal becoming and is also based on ideas from Fine (2005) (see section 2.5.2). I will not give an in depth analysis of this view but will note in places where the views contain similarities and where they come apart. In addition to these there are many different attempts to modify the standard accounts to resolve their difficulties. There is not space to lay out all of these but I will consider some particularly pertinent ones.

possibility of flow with respect to hypersurfaces and notes that some solutions to general relativity do admit a single global time function that is naturally privileged. These come from, for example, matter taking the general form of a perfect fluid with simultaneity hypersurfaces that can be defined orthogonal to its four-velocity. But these solutions are highly specific and we have no reason to favour them over the many other possible solutions that don't have this convenient global time. Even within this group of solutions there are many different options for what the general form of matter is and sometimes it leads to two or more privileged foliations rather than the single unique option we are looking for. Dieks (2006) also points out that the solutions that admit unique global time functions often involve generalisations and approximations of properties. The Robertson-Walker metric is a common example of this that assumes homogeneity and isotropy, which do not hold on a local scale. On top of this Dieks highlights that we do not have any experience of a universal cosmic time; the time between any two events is determined by the proper time along a worldline and makes no reference to this privileged time. There is no clear connection between time we experience locally and the privileged spacetime foliation. Overall the search for a universally defined time has very few feasible candidates and the possibilities that do exist have little to recommend them.[4]

Taking this further it is clear that on top of the search being fruitless it also lacks a clear motivation. Most of our physical laws are local and do not require universal foliations. The driving force behind attempts to reinsert it into our theory comes from our own experience of time and the intuition that the moment we experience is shared by everyone else in some way. But this falls apart under even a little examination. Firstly the effects of relativity at the scale we operate on are too small to make any discernible difference to our everyday perception of time and so we essentially do experience the same flow without it needing to be enforced by a universal time. And secondly our experience consists of localised happenings (or is at least limited to a finite region around us). There is no reason to suppose that our local experience is shared at other locations.

All in all the best way to reconcile flow with relativity is to abandon the requirement that it is universal and allow for localised models. Localised Temporal Becoming

---

[4] See Earman (2008) and Dieks (2006) for more details on this.

avoids this problem by using becoming that occurs along a worldline. Moments are replaced with spatiotemporally localised *events* (see Arthur 2019).[5]

Localised temporal becoming is not the only model of time to accommodate special relativity. Presentism is commonly objected to on these grounds but Brading (2013) suggests defining the present - and presentism - in terms of local dynamical laws instead. Earman localises the growing block model by similarly considering becoming along a worldline (section 2.3 will address this). Skow (2009) attempts to do something similar for the moving spotlight view with the same method of replacing instants with spacetime points.

2.3.2 The Privileged Present

Our human experience gives us privileged access to only the present time that we are right now experiencing. This leads to the view that is described as the standard view of the present (this label comes from Fine (2005)) where there is a unique privileged present moment and absolute facts about time and reality are orientated with respect to it. However, there are a number of questions arising about what a present moment is and what picks a particular moment out as privileged. These questions ultimately motivate localised temporal becoming's abandonment of any privileged present and the commitment to treating all times as equal (what Pooley, following Fine (2005), calls the non-standard view).

There is no indication of a privileged present in any physical theory; theories either describe the evolution of the universe as a whole or are limited to specific processes that could take place at any time. *Now* is treated no different to *here*; the spatiotemporal position we occupy is no more remarkable than any other. Here is only special because it is the basis from which we experience the world and if we were at any other point that would seem equally special to us. There is no reason that this should not also be the case for time: *now* is the moment we are in and hence it has precedence over any other moment for us, but this would be true if we were at any other point. In addition to the difficulties of identifying which moment is privileged it also seems integral to the idea of flow that this privilege

---

[5] This feature, along with the different ways to define an extended present that the next section will consider, will be essential to the arguments of the next chapter.

somehow shifts and moments take on being the privileged present in succession. Presentism attempts to do this by saying that the present is all that exists and the facts about succession are contained within the present by statements of similar form to ‘x is not now true but it *will* be true’ (where x is a tensed event). This notion of passage is laid out by Prior (1968). But this is a weak account of passage; as Deng (2017) puts it “[w]hen we imagine time’s passing and think about it in this intuitive, A-theoretic way, the tensed facts from more than one time enter into the imaginative episode, because as soon as time passes ‘on’, another set obtains” (p 1125). Time’s flow seems to carry more weight than the simple fact that the present contains facts that *will* obtain. The moving spotlight is the other main view that relies on a unique privileged present (basic formulations of the growing block also include one but later adaptations drop the claim that there is a single unique present) and explanations of how the present ‘moves’ through time quickly fall into the difficulties of dual times that we will see in the next section.

An additional problem arises. As we have already seen the attempt to define a single universal moment is fraught with complications and pursuing localised versions of flow is our best option. But a privileged present moment cannot be localised spatially the way that flow as a whole potentially can be; doing so would somehow pick out one arbitrary area of space as having unique features. There is even less support for picking out a single *spatio*temporal location than there is for picking out a single temporal one. In addition, presuming that if this privileged present exists then it is integral to the idea of passage, we would be limiting flow to one area of the universe while the rest is static.

A final problem with defining the present is that the way it is understood in its everyday use is very different to the concept of an instant in science. The *specious present*, an idea presented by William James (1890) and now widely used, refers to the duration that we hold in our immediate consciousness as ‘present’. Contrastingly, the present in physics is an infinitesimal moment or a single point in spacetime, Robb (1914) described this as the *punctual present* and it is also called the *point present*. Which of these notions would be the best candidate for a privileged present? Physics seems to favour the punctual present. However, this does not do justice to the intuitions behind a privileged present that are rooted in human experience. Presentism seems to lean towards the specious present, if we take seriously the examples given by Prior (1968) of the death of Queen Anne and falling out of a punt. These are certainly not point events. But what then defines the duration of the specious present? It seems to vary depending on the events being considered and this is not a strong basis for a rigorous and scientific account of flow. We are also left with the question of what meaning to give the succession of instantaneous punctual presents *within* the duration of the specious present.

The alternative to the standard view of time, which has a privileged present, is the non-standard view in which every moment is equal. This is not quite the same as the B theory of time as it does not necessarily reduce reality to tenseless statements about the universe. Nor is it quite the same as eternalism that argues that all events exist statically in a block. The non-standard view leaves open the possibilities that tense and existence are relative; each moment in time may give different facts of the matter and all of these different perspectives are equally fundamental.

The non-standard view does not abandon the claim that there is a present, only the claim that there is a unique present. Because the non-standard view does not require you to single out just one moment it allows for a much broader range of interpretations of what the present is. Spacetime is constituted by an array of punctual presents. All of these punctual presents are equal and you can define extended conceptions of the present with this basis without worrying too much about what this means for the internal structure of a present with non-zero duration. An example of this would be the Alexandrov Present as discussed by Arthur (2019) as a solution to this problem for localised temporal becoming that uses the idea of the Alexandrov interval, also called a causal diamond. This is defined as the intersection of the future light cone and the past light cone of two timelike related events. This allows for the present to be defined in terms of processes with a finite duration that one might encounter in the everyday world. Two point events such as the beginning and end of an extended event (such as the death of Queen Anne that Prior uses as an example) define an extended present. This provides a more rigorous scientific definition to the sort of concept described by the specious present. Another alternative is something like Brading's (2013) proposal to define the present in terms of what is singled out by a set of dynamical laws. This proposal has many interesting features to it but I will not consider it in depth here. The next chapter of this thesis will rely on ideas such as the possibility of an extended present but the precise details of how to define it are not relevant.

2.3.3 Dual Times

A recurring issue that pops up when trying to account for flow is that of doubly representing time; this problem can also be framed as vicious circularity or an infinite regress of time dimensions. This has been one of the strongest arguments in favour of adopting a static eternalist block model of time and has strongly motivated the particular form of becoming that localised temporal becoming argues for.

The basic problem is as follows: if we take flow to mean that something about time is moving then it must move with respect to something. Movement is defined as change with respect to *time* so to posit that time moves requires either change with respect to itself (circularity) or change with respect to second dimension of time (regress). If we take the latter and analyse the second time dimension the way we did the first then it requires a third level and so on, resulting in infinite regress. This argument, laid out clearly by McTaggart (1908), has become a staple objection to the passage of time and appears in many different variations. The moving spotlight view is particularly affected by it as it relies on a moving present. The problem, However, goes deeper than just a consideration of movement. The core problem is that time is doubly represented: once as moments in the physical world and secondly as the time which flows. The growing block exemplifies this as it has layers of time being added to reality and hence includes both the time represented by the layers and the time in which these layers are successively added on. This can be interpreted as movement – though as we shall see Earman tries to reformulate the view to avoid this – but more fundamentally it seems that the moments that are being added *are* time and should not be treated as something that is created or effected *by* time. In other words we should not need another set of moments to explain how the original moments are behaving because this ignores the nature of what moments are. How localised temporal becoming deals with this problem will be looked at in more detail in section 2.5.

2.3.4 Tensed versus Tenseless Existence

The final problem to consider is the distinction between tensed and tenseless existence that forms a foundation of much of the discussion about time and time's flow. It is central to the arguments that the rest of this chapter will make. It is not straightforwardly a problem as opposed to a means by which to articulate different positions. The distinction between tensed and tenseless *facts* is the basis for the A/B debate. Tensed facts are ones that hold now; in other words the fact is indexed to the time in which the statement of fact is made. The A-theorist position is that tensed facts are essential to our understanding of time and are non-eliminable. On the other hand, B-theorists of time take tensed facts to be reducible to tenseless statements where 'now' is replaced with the specified time: 'x at 11pm on Monday the 7th of December' (where x is a placeholder for some statement of fact, for example 'event e occurs' or ' object O exists'). These two positions separate out a dynamic view of time - in which tensed facts successively obtain - from a static, tenseless, one with no temporal flow.

This distinction becomes a problem, rather than just a principle of a particular position, when applied to the adjoining debate about existence. Presentism and eternalism are frequently taken to correspond to A-theoretic and B-theoretic positions respectively. (Essentially any model of existence that posits that time flows and there is a difference between present, past and future is taken to be A-theoretic, including the moving spotlight and growing block views; these views contain a commitment to the irreducibility of tensed facts.) However, these positions are primarily about *existence* and not *facts*, which is what the distinction originally applied to. Both presentism and eternalism make claims about existence simpliciter and do not commonly make a distinction between types of existence. Presentism states that only the present exists while eternalism claims that all of past, present, and future exist equally (and does away with these categories entirely). The growing block view similarly claims that past and present exist but the future does not.

Applying the tensed/tenseless distinction to existence is straightforward. Tensed existence describes events that exist *now* whilst tenseless existence is a more general sense in which an event *has, does* or *will* exist at some specified time. Dorato (2006) uses this idea to argue that presentism is either trivial or incoherent depending on how the main tenet of presentism – that the past and future do not *exist* – is interpreted. If 'exists' is understood in the tensed sense, then presentism is arguing simply that the past and the future do not exist *now* in the present moment; this is a trivial conclusion that no one would disagree with. But if 'exists' is taken as tenseless then saying that the past and future do not exist means that they *cannot ever* exist, for to exist at any point in time means they must have tenseless existence. Since the presentist wishes to retain the passage of the present from past to future it is incoherent to deny their existence so completely. This argument can equally be applied to basic formulations of the growing block that state that the future does not exist. Eternalism denies that that there is any sense of existence other than tenseless.

The problem of tensed and tenseless existence is also strongly connected to the dual times problem. The problem can be formulated as follows. At a certain time an event e_1 has tensed existence and all other events e_k have tenseless existence. But from any other time e_1 has only tenseless existence and the event at that other time is the one with tensed existence. Since all events should be treated equally (under the non-standard view) every event seems to have dual existence. This corresponds to the standard formulation of the dual times problem. There is the time series which each event tenselessly exists in and then another time series that singles out which moment is the present (and confers tensed existence on it).

This distinction is central to localised temporal becoming and is the starting point for how its ontology should be understood. However, the distinction has also been used in the growing block theory - as the next section will look at.

2.4 Against Reformulations of the Growing Block

The growing block is where the notion of temporal becoming (in its original conception) has received the most attention, and where its association with existence is most thoroughly explored. However, the dual times problem, along with the distinction between tensed and tenseless existence, presents a serious obstacle to it. The conception of becoming as coming into existence must be able to avoid the idea that time is both the moments coming into existence and the times at which these moments are added to reality. Before moving onto a focused examination of localised temporal becoming this section will consider one final view that claims to avoid all of the aforementioned problems. This is Earman's (2008) reformulation of the growing block view, which is articulated as a response to the deflationary B-theoretic version of temporal becoming as proposed by Dieks, Dorato, and Savitt. He takes this reformulation to be the most defensible version of temporal becoming that has sufficient metaphysical weight to be worth considering.[6] Considering this, and showing why it ultimately fails, makes it necessary to find a new conception of the connection between becoming and existence.

Earman presents a version of the growing block that is (he claims) compatible with relativity, has no moving parts and has a very different account of existence.[7] His model is actually a set of spacetime models, each one taking the form of a block universe but differing infinitesimally in size. These blocks are paired with an ordering relation, $\precsim$, defined between two models **n** and **n'** that are part of the total set such that **n** $\precsim$ **n'** iff **n** can be isomorphically embedded as a submodel of **n'**. In other words **n** comes *before* **n'** if **n'** is embedded **n**. This can be extended to relativistic spacetimes as well. Earman considers both hypersurface becoming and worldline becoming; the former can be discarded as Earman points out a number of problems with it and because, due to our consideration of localised flow, the search for universal hypersurfaces does not seem necessary. Worldline becoming requires simply that we replace the set of models of Newtonian spacetime with a set of worldlines, γ, in a relativistic spacetime, $\mathbf{R}$.

Supposedly this avoids the common problems given above. Earman avoids falling into the trap of tensed versus tenseless existence of the future by abandoning altogether the claim that the future does not exist.

---

[6] Another more recent defence of temporal becoming and the growing block can be found in Grandjean (2022), which attempts to rehabilitate the growing block and the open future in the context of theories of quantum gravity such as causal set theory. I will not consider this here as it does not relate to the more emergentist line of thought that I will be pursuing or the more specific questions about existence discussed here.

[7] It should be noted that Earman does discuss various problems with this view as well and does not advocate it outright. His argument is that it is the most defensible notion of becoming.

> “The advocates of the growing block model should concede that they cannot coherently sustain an absolute denial of the reality of the future; that would require giving up the growth in the growing block model in favor of some one fixed future-truncated block model. For the growing block model the denial of the reality of the future can only be perspectival, e.g. from the perspective of the world n_{2008} as it has grown up to T = 2008, future times are unreal but this denial of the reality of the future amounts to no more than the triviality that the future does not exist at the present time” (p. 144)

So Earman accepts the issues raised by the tensed versus tenseless distinction, as well as the dual times problem, and instead has larger blocks containing future events exist entirely separately. At one point he refers to them as being in “every relevant sense different possible worlds” (p. 143). The problem of tensed/tenseless existence was that denial of the existence of the future either meant that events do not exist *now* (trivial) or that events do not exist at any time at all (unsupportable). We now have that later events do not exist now but they do exist in different worlds that are larger than our own current one. But it seems that in doing so he has undermined the notion of existence itself. Each **n’** contains, by the definition of the ordering relation, a perfect copy of all the events existing in the smaller **n** of the previous spacetimes, so if all of these spacetimes exist equally then we have multiple copies of the universe. And each moment is represented many times over in each of the different blocks. It is unclear how this conception of existence can be applied to any of the familiar worries. Our tensed claims about the future are now taken to apply to different possible worlds as opposed to the one we are in. This seems to be a different question entirely; one not about perspectives but a multiverse of separate universes. Deng (2017) shares this worry:

> “But then further questions seem to arise at this point about how we should understand the perspectivalism of these truths, and the ontological status of each spacetime model in **R**. As Earman stresses, perspectival truth is not to be understood as on the block model, signalling that relative to each time, later times don’t yet exist. But how then? The sense of perspectivalism in question is starting to look less familiar

> than it seemed at first sight. Moreover, it's clearly central to making sense of the view under discussion. When it comes to this task, Earman's formulation would appear to be of no more help than the more familiar 'sequence of inked columns of increasing height' (Earman 2008, 138)." (Deng, 2017, p. 1116)

Even if one can reconcile oneself with this sort of existence there are potentially even deeper problems with this sort of perspectivalism. Earman claims that no extra dimension of time is required (avoiding the dual times problem) because each block in itself does not grow, instead the next moment is a completely new block. One could deny outright that Earman has avoided the dual representation of time at all because firstly time is presented *inside* each model as the set of moments that that model consists of and secondly time is presented as the set of the models themselves each of which represents a new moment. Earman would reply that this is a misinterpretation because each of the set of models is not tied to moments of time but are simply ordered by the relation, $\precsim$, which is based on whether the model is a subset of others and is both antisymmetric and connected. Time itself is still contained entirely within the blocks and the order of the blocks merely represents facts about what the process of becoming does to the sum total of reality. It may seem an unintuitive sort of flow but Earman acknowledges this and his defence is that "animation belongs to the province of Disney" (p. 140); what he is aiming for is to explain becoming and the different stages of reality that make up the passage of time.

The ordering mechanism seems a dubious replacement for a time series, but it does do the job required. A basic worry would be that is is simply disguising the problem. We still need something to associate a block that consists of times $t_1 .... t_k$ to reality at time t_k and not a time t_n where $n < k$. From the time t_k this block contains all the times that are present or past for this time; the block accurately represents reality as of time t_k and only time t_k. To do this Earman relies on Broad's original definition of the present moment as the moment that precedes no other moment. This means that a moment can only ever be called present if it is the last moment in the block. So the block $t_1 .... t_k$ itself singles out time t_k as the present by virtue of it being the final moment. This definition of the present transforms the ordering relation into a de facto time series; the blocks are arranged in an order such that it forms a chain of successive present moments. But, as it is not fundamentally posited and instead derived, it avoids the dual representation problem.

But what has happened to the notion of becoming in this picture? Becoming was proposed as adding a slice to reality and actualising a previously non-existent future. It now reduces to nothing more than the ordering relation that links one block to the next in a multiverse of different sized worlds. There is no longer any sense of transition, nothing crosses from one block to the other. They simply stand in an ordered relation to one another, and we can make little sense of how to make claims about other blocks. This is no less deflationary than the B-theoretic versions of temporal becoming that Earman criticises. It also, to return to the problem of existence and the status of the future, gives no meaningful sense in which the future comes into being. Future moments are simply alternate universes, each one completely static. To accept this view requires us to accept a radically different and largely unmotivated concept of existence in a multiverse and does not capture anything dynamic that becoming was supposed to represent.

2.5 Existence and Time

I turn now to consider localised temporal becoming and how it can be connected with a perspectival characterisation of the relationship between presentism and eternalism. With this in mind, I will argue that this account is not a tenseless block universe with a thin notion of becoming added at each moment, as critics have claimed. Instead, it presents a novel way of understanding how moments existing in a tensed manner are woven together by becoming, and the tenseless block is what this looks like from the outside. The resulting picture contains both tensed and tenseless existence but a focus on the latter obscures the dynamic importance of the former. This focus on the view as being tenseless (as evidenced in how it has commonly also been called tenseless passage or temporal B-coming) was meant to emphasise the benefits of the view in how it conforms with modern physics and rejects the idea of a privileged present. But ignoring the tensed aspect is a mistake and the tenseless block proposed by this view is radically different to the basic eternalist block universe.

2.5.1 Tensed and Tenseless as Perspectives

Savitt (2006), and similarly Rovelli (2019), give a careful account of existence and tense and propose a new way to understand the relationship between presentism (tensed existence) and eternalism (tenseless existence). Savitt argues that reducing to either tensed or tenseless statements always runs into difficulty. If there are only tensed statements then we fall foul of the problems Dorato gave with presentism. But if only tenseless statements are true then one cannot adequately explain how existence at one time is connected to existence at another. Savitt (2006) argues that this is because the difference between these views is not a fundamental difference but just one of perspective. Tensed statements about existence are ones made from *within* time while tenseless ones are made from *outside* of it. Existence can be viewed from either perspective and there is no straightforward notion of existing that can be made sense of without specifying the perspective.[8] To ask whether or not something exists simpliciter is incoherent; the question must be whether it exists from this perspective.[9]

---

[8] Savitt explores this in much greater depth. It is also worth noting that different authors use different definitions of the word tenseless. The use here is that corresponding to Dorato. Savitt replaces it with 'detensed' and uses tenseless interchangeably with atemporal. Arthur, who uses this distinction as well, follows Dorato with the use of tenseless and atemporal as separate notions. Savitt also switches freely between A versus B theories and presentism versus eternalism as they are concerned with tense and existence respectively and we are here dealing with the intersection of these notions

[9] Torrengo and Iaquinto (2019)(see also Lipman (2015) and Torrengo and Iaquinto (2020)) propose a somewhat similar idea (this also draws from ideas on perspectives on time from Lowe (1987)). They develop a view called Flow Fragmentalism (see footnote 2). They propose that there are *tensed* facts about all times that constitute absolute reality - these are not perspectival and there are no tenseless facts - and also facts that *obtain* locally. The set of absolute facts can be split into many fragments and need not be coherent across fragments (e.g. one fragment contains 'Socrates is now sitting' and another 'Socrates in now standing'). Flow is identified as the change in which facts obtain from fragment to fragment, but all these facts are taken to absolutely constitute reality and are not to be understood as relative. This is similar to localised temporal becoming except that it largely ignores existence in favour of facts and it denies that the absolute set of facts are in themselves a perspective comparable to the local perspectives (which are viewed as derivative from the absolute set of facts). This distinction between the absolute set and locally obtaining facts is not equivalent to the perspectival view of Savitt as the absolute set of facts is tensed rather than tenseless. However, in many respects the distinction between a global set of facts, and facts that obtain locally, is similar. Much of this chapter will be arguing that eternalism is best viewed as one perspective among many and that a single absolute set of facts should not be prioritised; this is essential to understanding why flow is an intrinsic

The two sorts of existence are not separate or independent from one another. If they were this would lead to the dual times problem as discussed in section 2.3.3 and 2.3.4. Once one switches to understanding tensed existence as a perspective inside of time, this issue dissolves. Tenseless and tensed existence can never come into conflict as they are simply different perspectives on the same thing. Positioned at a certain time all that exists in the tensed present moment; no other moments exist (in the tensed sense). But this is true for any moment not just of a single privileged present. However, when looking from outside of time, and not indexed to any specific moment, all moments exist in the tenseless sense. Existence is fully perspective dependent.

Put in these terms neither position can deny the other side of the distinction; the disagreement is instead about whether one is more fundamental. Fundamental could mean a variety of things. Clearly it cannot mean *ontologically* fundamental; after all these are different perspectives on the same ontology of events. There is nothing there to indicate that one perspective is more fundamental than the other. There are, of course, broader options for metaphysical fundamentality than just ontological.[10] I will not explore all of these options here; the aim in this chapter is to defend a version of becoming that is closely connected to existence (the next section explicitly lays out how localised temporal becoming uses tensed existence, although much is preempted here), and the ontological equivalence between the two perspectives is enough for this. The next chapter will expand on the idea of perspectives and connect it to the information available in any given perspective. This opens up other fundamentality relations, in particular the idea that the external perspective has maximal information and the embedded perspectives have limited information - this suggests the internal perspectives could be derivable from

---

part of reality. Flow fragmentalism rejects this and as such, by the arguments presented here, its account of flow is weak. Additionally as flow fragmentalism is largely formulated in terms of facts, it is much harder to connect to physics without developing a detailed theory of how facts are grounded. The primary aim of this thesis is to look at the physical basis for a perspectival view of time so this would present a major challenge to considering flow fragmentalism in depth here and I will not attempt to do so. The next chapter will point out places where localised temporal becoming could be given a more detailed theory of how to ground truth claims, but this would be a separate project.

[10] For example grounding or supervenience relations, or Sider's (2011) concept fundamentality.

or grounded on the perspective with maximal information.[11] Although, another way of understanding the relation between perspectives would be that the *total set* of internal perspectives cover all possible information - including information about the relations between moments (such as probabilities and uncertainty about the future) that is omitted from the external perspective; in this view the external perspective would be an abstraction from the total information contained in the set of internal perspectives and the fundamentality relation would go the other way. I will leave these options for exploring fundamentality open and instead focus on the more conservative claim that at the very least the internal perspectives are indispensable and that they are primary when it comes to understanding flow. The next chapter will also connect the internal perspectives to emergence and look at how they provide novel explanatory value; hence, even if there is derivability between perspectives, we cannot eliminate the internal perspectives and would instead have an emergent relation between them and the external perspective. Fundamentality would then rest on whether the emergence is strong or weak (where weak implies compatibility with reduction - the thermodynamics literature that I use in the next chapter generally assumes weak emergence). But this is hard to apply to the thinner notion of perspectives used in this chapter, which just focuses on existence and whether we are viewing it from inside or outside of time. For this ontological question, the perspectives are on a par.[12]

Borrowing from the next chapter, the most pertinent sense of fundamental for our purposes here is *explanatory* fundamentality. (By which I mean nothing more than something being necessary for a complete explanation of reality and irreducible to any other feature). Hence the question is this: which perspective should we look to for an explanation of the flow of time?

The basic eternalist view assumes that the external perspective should take priority and the internal perspectives can be derived from it, as a result allowing it to

---

[11] There are multiple ways to spell out what this amounts to metaphysically, see previous footnote.

[12] This particularly means that while becoming is essential to understanding the block at all levels (as the rest of this chapter will show), the features which make becoming a substantial account of the flow of time (the focus of section 2.6 and the next chapter) are emergent features. This would make flow an emergent phenomena even if a thin notion of becoming is essential to the block at all levels.

explain everything the internal perspectives can. However, this is an assumption that is often undefended and taken for granted.[13] If asked to give support the main answer is that this view is in line with physics. Physics is often seen as taking an objective external perspective that tells us how the world is without any reference to observers, and the block universe reflects this. The block also does not run into problems with special relativity and how to define a present moment (although as I have shown there are ways to overcome this in the presentist perspective). Theories like the differential geometry formulation of spacetime in general relativity also seem a good fit for a block universe. But, as I laid out in the introduction, this conception of physics as necessarily taking an external perspective is misleading. For example, Wallace (2019) argues that it is a mistake to eschew coordinate-based approaches to physics in preference of differential-geometric ones (which are a natural analogue to internal versus external perspectives on spacetime). Working in the former can help resolve problems and reveal structure that is hard to get at from the latter. Solving any practical physical problems requires adopting a coordinate system. The next chapter will give the arguments from physics more attention. For now I will focus on the metaphysical justifications for the internal perspective.

When considering an event coming into existence the tensed perspective is far more important. This is, in essence, the perspective that any physical object will itself take (in so far as we take physical reality to be constituted by objects existing in space and time and putting aside questions of fundamental ontology and its relation with spacetime). The sort of existence we are interested in is that of physical events and these events are closely tied to the constraints of spatiotemporal laws. The event from its own perspective occupies a tensed position and other events in the relation of *later than* are in the future, and events in the relation of *earlier than* are in the past. The external perspective is not one anything could ever occupy; even if we can choose to imagine an event as if from an external perspective. This is why the external perspective or ‘view from nowhere’ is also often referred to as

---

[13] Although A-theorists of time *have* argued at length for the indispensability of tense; I direct this comment particularly at the eternalism vs presentism debate and the assumption that an external perspective is the default position from which to approach both physics and metaphysics.

a 'god's eye perspective'.[14]

That nothing can ever occupy such a perspective does not absolutely preclude it from being more fundamental than the internal perspectives. This does not, for instance, affect the ideas about the external perspective having maximal information that are discussed above.[15] But it is an important consideration for the project of explaining features of time such as time's flow. Going back to the introduction, our primary evidence for flow is that it is a feature of our own experience, which is embedded within time. The concept of becoming is also explicitly defined (here at least) as coming into existence; this does indeed make little sense when considered from a tenseless perspective outside of time, as the critics of becoming have identified. When taking a tenseless viewpoint we should not necessarily expect to observe flow. This is exactly what the B-theorist does: they remove themselves from the spacetime block and see it as static and unchanging. The tenseless block may be a very useful way to think about time in some instances but there is a tendency to carry over temporal features of our embedded experience to this external viewpoint. One of these is the expectation to find dynamic flow when looking at the tenseless block even though tense - and changes in tense we might label flow - is contained *within* the block. To be able to see changes in tense from outside the block we would have to be able to observe one thing at one moment and another thing at the next; an observer experiencing 'moments' implies they are within time. This is why problems of double representation so often crop up for models of flow.

Both the internal and external perspectives are necessary for understanding time. But an eternalist assumes without justification that the internal perspective is merely derivative and the external perspective takes priority. (Likewise the presentist mistakenly assumes they can do without the external perspective and runs into

---

[14] We might question whether the external perspective is really a perspective at all as opposed to a lack of perspective. I will refer to it as a perspective here as it emphasises its connection to the internal perspectives. I take there to be no real weight to this question.

[15] Although, it may have bearing on other ways of assessing metaphysical fundamentality. Some approaches to understanding how science tells us about the world (for example Chang 2022; Barad 2007; Massimi 2022) take a pragmatic approach and base the development of scientific concepts around empirical access to the world. In these cases, the fact that we can only make indirect inferences about the external perspective may have weight as an argument against it's fundamentality as a concept for understanding the world. Structural realists, such as Wallace (2013) or Ladyman and Ross (2007), are unlikely to be swayed by this.

difficulties accordingly.) The tenseless block proposed in conjunction with localised temporal becoming is radically different from this. The correct way to understand this version of a tenseless block is as just one perspective among many. And when it comes to explaining features of time such as flow, the internal presentist perspectives should have primacy. It is a futile project to attempt to explain a concept such as flow, which clearly relates to the embedded experience, from an external perspective.[16]

2.5.2 Localised Temporal Becoming

To see exactly why the presentist perspective is more explanatory when it comes to flow I turn now to how localised temporal becoming uses it in conjunction with the concept of becoming; much of this was already preempted in the previous section, but it is important to spell it out thoroughly to avoid confusion. As section 2.2 mentioned, a perspectival view of the ontology is not an explicit commitment for all of the authors proposing this view, and a few directly oppose an ontological reading of becoming. But for many this position is compatible with what they have presented and helps to articulate the view in stronger terms. With this ontology localised temporal becoming is capable of delivering both aspects of what is required from an account of flow. To reiterate, these were: a) it captures the sense that reality changes from moment to moment and b) it describes what the relationship between moments is (i.e. how one moment is transformed into the next).

What this view argues is that temporal passage is the succession of tensed becoming. Becoming refers to coming into existence in the tensed sense. Because becoming is necessarily tensed there is no sense in which an event *already* exists at times before its own time; at each moment a new event comes into being and there is a change in what exists. Hence the coming into being of events throughout time constitutes the flow of time.

Where this view is commonly misinterpreted, as Pooley and Earman do in their criticisms, is in interpreting the tenseless block in this account as the basic eternalist

---

[16] One might worry that a eternalist could reply that flow is not a feature that *needs* explanation. The work of the next chapter will show that it is an important part of understanding how physics makes predictions and retrodictions.

block universe. Instead the last section showed that the tenseless block used in localised temporal becoming is radically different and should be understood as just one of many perspectives on time. Becoming is not a trivial afterthought but an essential way of relating the different perspectives; there are many internal tensed perspectives, one for each moment (spatiotemporally localised events), which are woven together by becoming. The tenseless perspective is a way to view the block of time from the outside. The internal perspective *starts* with the tensed becoming of an event at its spatiotemporal location and becoming relates all the events to each other.[17]

Temporal flow must be understood as what ties together the internal perspectives. To further illustrate how to think of this we can make a comparison to the work of Kit Fine. Fine (2005) approaches the problem of tense with a different basis but it carries over well into this presentation as it is informed at least by the work of Stephen Savitt if not others working on these views. He talks about how philosophers have commonly misinterpreted the claims that realists about tensed features make:

> "For them, it is as if the realist has wanted to assert that each time or time-slice of the world is present but subject to the qualification that this should only hold at the time in question; and put this way, it is hard to see how it might amount to anything more than a triviality to which even the anti-realist could agree. But the proper formulation of the intended claim is that reality is constituted, at each time t, by the fact that t is present. This is quite different." (Fine 2005, p. 280)

The basic difference described, adjusted to fit the language of existence of events used here, is that in the first case existence is being considered from a bird's eye view position outside of the static block and we can easily see that it is a trivial fact that each event exists at the time it occurs. The static block represents absolute reality; this is the basic eternalist view. But in the second case, which corresponds

---

[17] Arthur (2019) notes that if the block universe can be said to exist at all "over and above the events in it" then all it is is "the metrical, topological and ordering relations among events" (Arthur, 2019, p. 55). This gives all the necessary structure that goes beyond the tensed events within the block. The questions of how events in the block are given structure by laws and how this relates to the metaphysics of time laid out here go beyond the scope of this thesis.

to the block used in localised temporal becoming, we are reinstating ourselves inside of the block and considering reality *at that time*. What exists as of a point in time is truly relative because it will be different depending on which time you consider it from, but at each moment there is a definite (tensed) fact about what exists. This difference in the very makeup of existence from moment to moment constitutes a richly fundamental view of change and passage throughout time.

We can now evaluate this against our two requirements for flow. The first - change in reality - is captured because becoming is a relation between two tensed realities. What exists from moment to moment changes in the internal perspective. The second aspect was how one moment relates to the next and the sense of flow as a dynamic transformative process. In part this is satisfied by the definition of becoming as coming into existence; a previously non-existent moment is transformed into an existing one. But this doesn't seem to fully explain how one moment relates to the next. It is also vulnerable to the objection that we are defining becoming as occurring along a worldline, but nothing distinguishes timelike from spacelike worldlines. Becoming seems to be a relation between *any* spacetime events, not just between one moment and the next. Yet flow is specifically a temporal notion and should not occur between spacelike separated events. To answer this we need to look closer at how this model treats future moments and what happens when they come into (present) existence. Without this becoming may seem like just a structuring relation among moments rather than a truly dynamic process.

2.6 An Open Future

To see the value of the presentist perspective in providing dynamic flow I will turn to the idea of the open future. For many (Broad 1923; Pooley 2013 and many others) the actualisation of a non-existent *open* future is the key to understanding the flow of time. The difference between the past and future is that the future is uncertain and there are many possible ways the world could turn out to be. Why the open future is so important and what the details of this are is a larger question that I will explore in more depth in the following chapter alongside considering why, if the future is open, the past is not. Here I will focus on the question of existence and simply show that an open future is compatible with the ontology of localised temporal becoming. This claim is initially at odds with there being a tenseless block of all events so it requires careful articulation. The following chapter can then build on this basis to show how an open future connects becoming to the directionality found in physics and the way that physics treats the openness of the future.

Various models of the open future have been developed with a range of different ontologies. Some of these limit the open future to being purely epistemic (Lewis 1986; MacFarlane 2003 for criticism). Others (McCall 1994) take a strongly ontological approach where multiple futures exist in a branching future and branches 'fall off' as time passes. I will focus here on the more moderate view articulated by Pooley (2013), which follows from Belnap (1992; 2003; 2012), Placek and Belnap (2012).[18] The resulting ontological picture, I will show, has no meaningful differences to how localised temporal becoming represents time in terms of the perspectival switch between tensed and tenseless existence. Moreover the account of flow in the open futures model can be achieved by localised temporal becoming.

The commitment to open futures is not agreed upon by all the authors who have advocated for localised temporal becoming. Dorato (2006) denies that there is any sense of an open future; he takes the tenseless block to fix all relevant events. Others do take it to be a key consideration although they do not spell out in much detail how this follows from the ontological picture they present (for example Ismael 2016a). Most, however, do not mention it at all.

2.6.1 The Block Universe and Open Futures

The open future is not a statement about there being some sort of indeterministic laws. Rather the future is *indeterminate*. Most commonly this is taken to mean undetermined or unsettled truth values for future statements. As of a time t_1 facts about a later time t_2 are either true or false but it is not definite as to which. Facts are made true by the events at t_2 and prior to this time they are not fixed.[19]

---

[18] Pooley (2013) gives a critical analysis of the different models available and why this version of open futures is the most coherent picture for an A-theorist to take.

[19] This presentation allows for compatibility with bivalence, following Barnes & Cameron (2009) and Wilson (2013). Bivalence is maintained as future statements determinately have a truth value even if the state of the world at t_1 fails to determine which value it is. I will give more details of

There is a branching structure containing all possible futures; but this structure is not ontologically real - only one future actually becomes - it just represents these abstract possibilities. Pooley advocates for understanding this in terms of Stein's (1968) notion of relativistic becoming in which events in the past lightcone of event e_1 are determinate whilst everything outside is not and becoming then occurs point by point along a worldline. This makes clear why whether or not the laws are indeterministic is irrelevant to the future being undetermined. Events outside of the past lightcone can always have an unknown effect on the next event along the worldline so even deterministic laws would not allow the future to be predicted based on the determinate events inside the lightcone.

Pooley supports the non-standard view of time (no unique privileged present) and this means that there is no single 'tree' of time with a determinate trunk of the past and branching futures. Instead each moment in time is associated with such a 'tree' and there is a sequence of trees that corresponds to time as a whole. Initially this may strike one as being somewhat similar to Earman's set of spacetimes model adjusted so that the non-existent future of each block is replaced with branches of abstract possibilities. If this were so the model would fail on the same accounts. But Pooley is careful to clarify that he does not think that this sequence represents "how reality is absolutely" (p. 18) and this makes his view distinct from Earman's. Pooley articulates what he sees to be the correct way of interpreting the sequence in the following passage, which also shows how this view initially seems to differ from localised temporal becoming:

> "Just as the tensed facts that hold as of some time are not reducible to tenseless facts, there is no need for them to be deducible from the tensed facts that hold as of other times. . . .It is hard to see what the insistence that [tensed] facts are not reducible comes to, for there is a unique representation of reality—the block universe—from which the

---

these views on indeterminacy in the next chapter as they will be the basis for exploring the open future further. Pooley, who this chapter follows, rejects bivalence and instead views the openness of the future as a failure of bivalence (the truth values of statements about the future are *neither* true *nor* false). I have followed the former views as they provide resources useful for exploring the questions raised in the next chapters. However, the overall account does not rest on this and it could be reformulated to reject bivalence instead with no significant loss. The next chapter will discuss this further.

> perspectival facts can be derived. This is no longer true of the open-future model. The primordial branching-structure captures only how things might turn out, not how they will turn out. The block universe history that constitutes the ideal limit of the sequence of the model's branching structures not only does not correspond to the facts as of any time (the end of time is never reached), it also, when interpreted as representing the absolute facts, misrepresents as determinate future facts that are genuinely unsettled." (Pooley, 2013, p. 18-19)

This explores the problems Pooley sees with conceiving of a block whether it be fundamentally tensed or otherwise. There are two main claims here: first that, if a block can be conceived of, then it in some way fixes facts as determinate when they should not be, and secondly that there is no such block for the open future model. These taken together lead Pooley to reject a block universe and adopt an open future ontology. But if we adopt the perspectival view of the block universe it becomes clear that Pooley's first claim does not follow. A tenseless block, understood as an external perspective on the tensed events that occur from moment to moment, should not be understood to represent events as determinate prior to the time they take place. A block such as the one used in localised temporal becoming should not be mischaracterised as an absolute description of reality (from which all other facts are derived); it is just a perspective.

If the external perspective was ontologically more fundamental then this would be a basic B-theoretic version of the block and indeed offer nothing meaningful or interesting. Yet as we have seen there is no basis for prioritising the external perspective. It is true that from a tenseless perspective one can see all the facts of time as determined and static and the future does not seem open. But this does not imply that this can be carried over into a perspective inside of time; one cannot, as it were, jump out of time to learn the future and then jump back in. Facts are determined by the coming into tensed existence of events at exactly the time they occur and the tenseless perspective is a by-product. In the same way that events do not *already exist*, they are also not *already determined*. At a time t_1 tenseless statements about future events cannot be made because the nature of tenseless statements is that they are made outside of time. We can of course make statements of the same form but it would be a mistake to label them as genuinely tenseless. At the time t_1 we do have access to some determined events, namely all those in the causal past of that event (putting aside for now the question of why there is an asymmetry - this will be returned to in chapter 3), and so it is fairly simple to artificially imagine ourselves in a 'non-temporal' perspective to look at these determined events. But this is equivalent to the ideal limit sort of block that Pooley mentions; it is looking back from *after* time and not from outside it. Because we can make artificially tenseless statements about the past we suppose that we can make them about the future simply by imagining later times. But these are not genuinely tenseless statements (they are not *actually* made from outside of time) and do not mean that the future is closed. So the worry that tenseless statements fix as determinate facts that should be open is unfounded.

As for Pooley's second claim - that the open future model does not have a block - this can also be responded to by recognising the block as a feature of a perspective on time rather as an absolute entity. The open futures model is non-standard and all times are equal (each with an associated branching structure). Taking the tenseless perspective is to look at all these moments together; this does not imply a block that exists at some endpoint of time but rather it is the view from outside of time completely. From outside of time the branching structure, which is *not* part of the ontology, cannot be seen and what is left is a set of individual moments. This is a perspectival tenseless block exactly like the one in localised temporal becoming.

All in all the block universe of localised temporal becoming is the same as that of the open futures model when the perspectival view is adopted and it is distinct from B-theoretic blocks in both theories.

2.6.2 Passage

Now we can look at how each model achieves passage to see if there is any difference between them there. For localised temporal becoming this is the simply the tensed becoming of events. For open futures it is the closing off of possibilities by an event becoming determinate. The closing off feature of this sort of passage can only be resultant and not constitutive as these possibilities are not real and are not playing an active role in passage. Passage itself is the event becoming determinate.

Once again this is modelled on Stein's relativistic becoming. Stein discusses how real becoming can be formulated in accordance with special relativity in his 1968 paper as a reply to arguments from Putnam (1967) and Rietdijk (1966) that special relativity is incompatible with probabilism. He reiterates the points in 1991 against Maxwell (1985). Stein's argument is that the events that *have become determinate* as of a certain event e_1 are all – and only – those that lie in the past light cone of e_1; events in the future light cone and outside of the light cone of e_1 are indeterminate. What *has become* is spread over a past lightcone but becoming occurs only at the apex of it: at event e itself. Stein clarifies this point as a reply to Putnam who argues that between e_1 and an infinitesimally later event e_2 the lightcone increase in size across its whole surface and the new events on the new surface will "*have been* without its ever having been true that they *are*" (1968, p. 246)(emphasis in original). Stein points out that in special relativity "*an event's present is constituted by itself alone*" (p. 15) and therefore one event can have nothing to say about the present of another. This fits perfectly with the picture of localised flow along worldlines. When e_2 becomes determinate it is added onto all the possible timelike worldlines that can be defined in its past light cone including the new lightlike worldlines that make up the surface (which the lightcone of $e_1$did not contain).

This relativistic formulation of becoming is identical to that of localised temporal becoming: it occurs at a single temporal location and flow can be defined as the successive determination of events along a worldline. Pooley acknowledges the similarity between these accounts but claims the difference is that Stein does not believe in a block universe because the future is ontologically open. However, we have already argued that the block universe of localised temporal becoming is compatible with an open future.

Putting everything together the open futures view put forward by Pooley can be framed as having the same formulation of a block universe as localised temporal becoming and the same account of passage. Open futures focuses on the effects of flow and what this means for the existence of events while localised temporal becoming focuses on the mechanism of flow itself. Laying out the ontological picture behind localised temporal becoming, as this chapter has done, reveals their compatibility. Each chooses to highlight different features as what is most important. Focusing on the mechanism of becoming is a more direct answer to what the flow of time is, but this mechanism cannot be separated from how it relies on the tensed existence of events - and the open future that results from this - that the open futures model lays out more clearly.

2.7 Conclusion

The perspectival relationship between presentist (tensed) and eternalist (tenseless) views of existence reveals that localised temporal becoming gives an account of time's flow that is both coherent and meaningful. It captures the way that open futures are an essential part of understanding time and retains the benefits of the tenseless block. Understanding this view as fundamentally tenseless with becoming simply added at each moment is a misinterpretation; and accusations that it is trivial, which follow from this assumption, are likewise mistaken. The tensed perspective is equally, and frequently more, important.

Looking back at the introduction, I argued that perspectives, emergence and information would be the guiding principles that the metaphysics of time can be derived from. The focus in this chapter is on perspectives. Properly appreciating the perspectival basis of localised temporal becoming explains how this view accounts for flow. Localised temporal becoming provides both a) the change in reality from moment and b) the transformative relationship between one moment and the next. The relation between moments is how the future is undetermined with respect to the present and how flow transforms possibility into actuality.

The open futures model is a complete model of time in its own right, however, the perspectival language of localised temporal becoming gives this account more resources to work with than the open future model on its own. This thesis is primarily focused on the internal perspective but the external perspective also has many merits in certain situations where the internal relations of time are not relevant. Localised temporal becoming allows us to talk about the block universe in these cases, while the open futures model maintains that there is no block.

The physical basis for openness needs further exploration. The next chapter will explore this, and show how it leads us to ideas about information and emergence. While I have shown that the metaphysics of this model is coherent, there is more work to do in exploring the tensed view and in particular to argue that, despite the common (although by no means absolute) assumption that the tenseless perspective is more useful for the way we model time in physics, the perspective from inside of time is just as valuable. This will provide justification for claiming that flow is an ineliminable aspect of reality.

Specifically the next chapter will argue that from the tensed perspective the future is *metaphysically* indeterminate and not just epistemically so. Currently both localised temporal becoming and the open futures view are vulnerable to a response by an eternalist that claims that all facts about the open future can be derived from tenseless ones by just limiting the information available to a particular time and considering the range of futures compatible with that subset of information. In which case the tensed perspective appears redundant. There has so far been little justification offered either here or more widely in the open futures literature for what the eternalist is missing when they claim the the future is just epistemically open, other than our intuitive sense that claims about the open future have more weight that that. So I will look at what the metaphysical view does better, to provide justification for this model that is independent of just wanting to explain flow.

3 The Open Future and the Recorded Past

3.1 Introduction

Our experience of time is solely limited to the present, and yet the past and future play ineliminable roles. From the present we are able to conceive of the future in front of us and the past behind us. Typically, the future is seen as being open or undetermined in some sense and the past is fixed and determinate. This commitment is found frequently in the differing ways in which we talk in everyday language about the past and future. The flow of time is often associated with the closing off of future possibilities and transitioning to the fixity of the past. This chapter will explain what it means for the future to be open and what the physical basis for this is.

In the previous chapter I developed the view - localised temporal becoming - in which the flow of time is the coming into (tensed) existence of events along a worldline. The ontology behind this claims that presentism (tensed existence) and eternalism (tenseless existence) are two different perspectives; the former the perspective from inside of time (in fact this is many perspectives, each corresponding to a different moment) and the latter from outside of it. I also showed that this was compatible with having an open future. Events can be represented in the tenseless block and from this external perspective the future is fixed (in fact there is no separation between present, past and future). However, from inside of time there is no sense in which future events *already* exist at the present moment. The future can be described by a set of possibilities and as each new moment comes into being these possibilities are successively closed off. It is the process of becoming that weaves together the tensed moments into an eternalist block. The collection of tensed perspectives, and the open future they entail, form a complete description of reality. The external, eternalist, perspective is only one perspective among many.

The concept of the open future, however, deserves more attention. While I have shown it is compatible with - and indeed necessary for - the ontology of time and an account of flow, more work needs to be done to justify not only why this is important but also what exactly it means for the future to be open. Crucially, I will show that openness is part of the metaphysical facts about reality and is not merely epistemic (or rather that in many cases epistemic openness *is* metaphysical). Perspectival structure in the world - specifically the lightcone structure of spacetime as well as emergent theories such as thermodynamics which contain probabilities - leaves the future indeterminate. This will justify the conclusion of the previous chapter that, not only is the eternalist perspective just one perspective among many, but the internal perspectives are primary for explaining the flow of time and without them we do not have a complete description of reality.[1] The eternalist picture is leaving things out; it can reconstruct facts about the open future based on limiting what information is taken into account, but these come out as merely epistemic and it fails to capture the metaphysical implications of the open future. Perspectival openness treads a middle ground between the current prominent views in the literature regarding the open future. One side look for a strong sense of metaphysical indeterminacy in the fundamental state of the world (e.g. Mariani and Torrengo 2021a, 2021b; Barnes 2014) and dismiss more general accounts of metaphysical indeterminacy (e.g. Wilson 2013) as epistemic. The other side embrace the subjectivity of epistemicism and look for openness due to agency in the way we act and deliberate (e.g. Ismael 2023). I will show that metaphysical indeterminacy can be derived from perspectival - and not necessarily fundamental - features of the world and it is these features that produce the openness found in the agency account.

The perspectives I will be talking about are defined by physical structure and constraints that affect the way we gain and use information. I will often call this epistemic structure (and certainly they are often labelled as such in the wider literature), but it should not be dismissed because of its connection to epistemology; openness is epistemic but not *merely* epistemic. There are a multitude of ways in which the structures on which epistemic concerns are built are used in physics. I will show that what information can be discovered at different positions and at different scales is an essential feature of the world. As such, this view provides a

---

[1] As per the previous chapter, I leave open the question of whether the external perspective is more fundamental in some metaphysical sense (although the last chapter showed they are ontologically on a par) to focus on the question of explanatory fundamentality, and show that the internal perspectives are ineliminable.

metaphysical openness that is strongly rooted in physics.

Finally, the closing off of the open future *into a fixed past* is what makes becoming a good account for the dynamic nature of time; but the fixity of the past has received relatively little attention in comparison to the debate about open futures. To fully understand what it means to be open we must also have an account of what it means to *not* be open. Given the ontology I laid out in the previous chapter it is of vital importance that once an event has come into being it then goes out of existence again; else the view would collapse to some sort of growing block. But this leaves the past and future with equal ontological status and any reasons for considering the future to be open should also be applied to the past. I argue that the answer to this follows from the metaphysical account of openness I will present. Specifically, I propose that the past is fixed because physical records of it exist in the present that can be used as a basis to make definite true or false claims about the past. This in itself may seem controversial as a metaphysical thesis and it leaves the past occasionally open (and the future occasionally fixed), but the claim follows from the considerations I will give of how epistemic structures constitute essential features of the world. What is more, records are an important part of the physics of time asymmetry and this view allows us to connect the metaphysical asymmetry between past and future to the many physical temporal arrows (such as thermodynamics, agency, and causation). This yields a pervasive time asymmetry that falls out naturally from physics and is not an intrinsically posited feature of becoming.

3.2 Metaphysical Openness

What does it mean for the future to be open? This is generally taken to mean that the truth values of statements about the future are indeterminate. Statements are either true or false but from the perspective of the current time there is no determinate answer as to which.[2] I will present two accounts of metaphysical

---

[2] Another way in which truth values can be indeterminate is a failure of bivalence - future statements are *neither* true *nor* false (Pooley 2013; MacFarlane 2003) - but many challenge this and show that bivalence can be maintained (Barnes & Cameron 2009; Wilson 2013 - details to come in section 3.2.1). I will allow the latter view, but I will not defend it extensively here (Torre (2013) and Barnes & Cameron (2009) discuss it in more detail). My reason for following these

indeterminacy that claim that future truths are indeterminate. Both are similar and together provide the resources to understand what it means for the future to be open. Additionally both these accounts can be made more precise using the ontology of localised temporal becoming in which, from the localised presentist perspective, the future does not exist.

3.2.1 What is Metaphysical Indeterminacy?

The first view comes from Barnes and Cameron (2009; 2011). They argue that we should distinguish between settled and unsettled truth values. Take the claim 'there will be a sea battle tomorrow', as of tomorrow this will be either determinately true or determinately false. However, as of today we can only say that it has a truth value - it is determinately *either* true *or* false - but it is not yet settled as to which. The events of tomorrow will settle it one way or the other. This view, they claim, is compatible with an eternalist view of time. The tenseless facts about the world settle the truth values (and it is based on this that bivalence is retained). Added to this is the relational notion of settledness, which holds relative to specific times. Relative to today the future is unsettled, but the tenseless facts about the world ensure that there *is* a truth value. Although they do not put it in these terms, we can view the relation of settledness as referring to perspectival tensed facts.[3]

'Unsettled' is a relation between the present and the future and needs some further explanation to make sense of. There are multiple possible futures compatible with the laws of nature and from the present there is no way to settle which one is the actual future. This argument relies on reconciling the determinacy of laws with metaphysical indeterminacy. (Indeterministic laws would also make the future unsettled in this way but would limit the openness to just those specific laws, which Barnes and Cameron wish to avoid.) There are a number of ways to defend

---

accounts is that, in them, indeterminacy can be linked to predictability in physics (as the next sections will show); this provides a more physical grounding for openness than the accounts that reject bivalence have offered.

[3] Other ways of thinking about bivalence could be applied here. For example, one could argue that bivalence is retained in the external perspective but fails in the internal perspectives, although this introduces an additional layer of complexity to address. I follow Barnes and Cameron's arguments for bivalence here for simplicity. Wilson, whose view is considered below, also supports bivalence, see Wilson (2016)

this claim, which all hinge upon questioning what is available in the present and past that can be used to determine the future using laws. Barnes and Cameron argue that a best system account of laws cannot be worked out without looking at the entirety of time including the future, hence based on just the present and past it is indeterminate what laws obtain in the world. Alternatively, if one does not favour a best system account of laws, they also claim that even if we fix the laws there may be metaphysical indeterminacy in the present state of the world that precludes deterministic laws from fixing the future. But putting these arguments aside there is also a simpler justification that can be given in the context of the arguments developed in the previous chapter. If we are working with a localised present, and becoming occurs along a worldline, then information included in the present and past is not enough to fix the future. Influences from outside of the lightcone would have to be taken into account.

On this account, it seems, the openness comes down essentially to predictability. The future is unsettled because the facts of the present (and the past) in conjunction with the relevant laws are not enough to single out a unique future.

Wilson (2013) provides a similar account of metaphysical indeterminacy but instead of 'settled' truth values she uses the idea of determinables and determinates. This provides a more precise way of spelling out what Barnes and Cameron were getting at with the notion of a truth being settled or unsettled.[4] Wilson argues that we are able to make determinable claims about the future. The claim that 'there will be a sea battle tomorrow' is the sort of claim that a certain state of affairs will make either true or false.[5] However, at the present time the relevant state of affairs has

---

[4] I have only presented these accounts briefly here and their full details differ in exactly what they commit to, the difference between them has come to be called meta-level (Barnes & Cameron, also Barnes & Williams 2011) versus object-level (Wilson). For a more nuanced look at the differences between them see Wilson (2016). The main difference is that for Barnes and Cameron is it indeterminate (based on the present state of affairs) as to which determinate future states of affairs will obtain, while for Wilson the present state of affairs is indeterminate (i.e. it contains determinables without determinates). I will use ideas from both accounts in this chapter; however, it should be acknowledged that spelling out a fully detailed argument of how metaphysical indeterminacy works and how my proposals in this chapter can be made precise in metaphysical terms would require a more careful analysis of how these fit together. Here I will largely use conceptual ideas from Barnes and Cameron but the practical machinery of Wilson. My aim is to examine how metaphysical indeterminacy can relate to the open future and the physical basis for this rather than to produce a detailed account of metaphysical indeterminacy in and of itself.

[5] State of affairs is a complex term in its own right and it has been argued that states of affairs

not yet obtained - there is no determinate for this determinable.[6] When a state of affairs obtains at a time it acts as a determinate truthmaker for the determinable.[7]

Where this account is let down is that no further backing is given to what it means for a state of affairs to obtain. A model of time is needed to make sense of this. In an eternalist block universe the states of affairs of all times obtain and hence there is no indeterminacy. However, in models of time where the future does not (yet) exist, there is room to claim that future states of affairs do not yet obtain. The ontology of localised temporal becoming provides this. Future events do not exist from the present perspective, although even in the present we can claim that they *will* exist. As such the future is determinable, but its determinates do not yet exist. Until the time at which the determinates come into existence the truth value of future statements will not be settled.[8]

However, as Barnes and Cameron consider, another way for states of affairs to settle future truth values is if the present state predicts the future state. Wilson's account of determinables rests on the idea that nothing in the present acts as a determinate for the future, so this is a notable omission. She makes no mention of the possibility that a current state of affairs could act as a determinate for what state of affairs *will* obtain in the future. Taking this into account Wilson's indeterminacy also rests on a failure of predictability.

Together these accounts provide a precise and useful way of defining what it means for the future to be open. They are also supported by the argument that the presentist and eternalist views are different perspectives on time. The presentist perspective helps makes sense of what it means for truths to be unsettled or for

---

cannot act as truthmakers (a role that must be played by *facts* (see Textor 2020)). I will follow Wilson's (and Barnes and Cameron's) usage of the term.

[6] This is also compatible with claims for which the determinate will *never* obtain. There may be a determinable claim about the future (i.e. there will be a sea battle in the future) where no future state of affairs ever settles it (the possibility of a future sea battle remains open indefinitely as there are always future times where it may be made true).

[7] Truthmakers are often understood in terms of grounding, i.e the physical state of the world grounds the truth or falsity of claims about it. This is distinct from the more general (and more controversial) concepts of metaphysical grounding that are sometimes used, for example, in debates about causation. See also footnote 3.

[8] It should be noted that this goes beyond Wilson's (and Barnes and Cameron's view) somewhat and frames it in terms of the perspectival ontology I am using.

future states of affairs not to obtain.

3.2.2 Why is the future metaphysically indeterminate?

What we take to be a genuine failure of predictability dictates what sort of metaphysical indeterminacy we find in the world. One way that metaphysical indeterminateness has been construed is as being due to our current position in time i.e. it is perspectival. We cannot know the future because *we are not there yet.* This is often dismissed as being merely epistemic and not being grounds to consider the future truly metaphysically open.[9] Simply not being at the relevant time does not constitute metaphysical indeterminateness. The accounts of the open future, such as Pooley (2013) that do mention perspectival structure such as the lightcone, do little to justify why this should count as genuine openness and it is often not explicitly pointed to as the *source* of openness. Dorato (2006), in his presentation of localised temporal becoming, dismisses the open future on exactly these grounds.[10] Similarly Grandjean (2022) rejects perspectival accounts of openness as only weakly explanatory. A strong failure of predictability is looked for instead. Mariani and Torrengo (2021a; 2021b) provide one example of this approach. They argue for genuine passage of time based on dynamical reduction models of quantum mechanics (specifically GRW). They use Wilson's machinery to do this but reject the idea that indeterminacy is due to our position in the present preventing us from knowing what future state of affairs obtains. Instead they look for metaphysical indeterminacy in the present state of the world. This puts indeterminacy into the state of the world itself and our practical failure to make predictions is derivative of that. And as they see it, the only place to find this is indeterminacy in the quantum state. This indeterminacy entails the indeterminateness of the future.[11] The quantum state could collapse, in the future, to a number of different states.

---

[9] This is largely what Wilson suggests although, as stated, she does not consider predictability. Barnes and Cameron are also perspectival, but to a lesser extent as some of their considerations about indeterminacy in states and laws lie closer to the strong metaphysical view.

[10] He does not mention any accounts of indeterminateness specifically but makes the general argument that our position in time is a mere relational fact and should not entail any sort of openness.

[11] This is similar to Barnes and Cameron's view, which hinges having either an indeterminate present state, indeterministic laws, or indeterminacy in what laws hold.

Their view is coherent and well presented but its downside is that it commits you to a particular interpretation of quantum mechanics (and in Mariani and Torrengo 2021b a particular ontology of GRW on top of that).[12] It also depends on pertinent problems with that interpretation - such as the tails problem - having a viable solution (as they note in 2021a). It is an open question as to whether the GRW interpretation is the correct one, but there are many reasons independent of one's views on the openness of the future why one might prefer or reject this view.

This strong view sees only indeterminacy in the fundamental state of the world as important.[13] Any other failures of predictability due to perspective are merely epistemic and amount to nothing more than ignorance of the future. This would also rule out the argument - mentioned earlier - where a failure of predictability comes from the constraints of a localised present.

But a large part of the argument of the previous chapter was to establish that perspective is important; additionally, there are many physical limits that perspectives place on our ability to predict the future. As such, taking this strong view and relying on fundamental quantum probabilities is unnecessary. There are numerous other places in which we make probabilistic predictions about the future that should not be dismissed because they may be perspectival.[14] These limits are part of the world as much as the quantum state is.

The introduction (as well as the previous chapter) has laid out some reasons to

---

[12] Mariani and Torrengo do not consider other interpretations of quantum mechanics. It is not clear that anything in their view precludes other interpretations of quantum mechanics that have probabilistic collapse of the wavefunction. However, the Everettian interpretation and Bohmian mechanics would be ruled out as not having genuine collapse. For example for an Everettian the future is perfectly predictable and it is only epistemically open as to which branch you will end up in (or if you will end up in all of them). And subjectivist accounts such as QBism would not have the necessary connection to ontology to make becoming a metaphysical fact.

[13] Putting aside concerns as to whether quantum mechanics is truly a fundamental theory, which has largely been superseded by quantum field theory anyway in order of fundamentality. The GRW interpretation does also consider scaling up to macroscopic domains as important, but want to base this on fundamental microscopic collapse.

[14] The possibility of non-fundamental indeterminacy has started to be debated in the literature, Barnes (2014) argues strongly that indeterminacy must be fundamental and Mariani (2022) responds. This is somewhat at odds with the insistence from Mariani and Torrengo that quantum mechanics is the only true source of indeterminacy that could be suitable for temporal becoming. However, all authors strongly reject any indeterminacy which could relate to knowledge or perspectives, it is this strong claim that I will be arguing against here.

take perspectival views of the world seriously. And the rest of this chapter will go into more detail about the way that physics uses epistemic constraints to show that a failure of predictability due to perspectival limits *does* have an essential role in our metaphysics (and in physics itself). However, it is worth considering here a general perspectival constraint that physics places on knowledge: special relativity and the lightcone structure of spacetime.[15] The biggest recent push-back against presentist views of time has been to take seriously the way that relativity disproves the idea that there is a single universal present moment. The response has been to develop localised accounts of the present (and of temporal flow). Along these lines the limits that the lightcone structure puts on the predictability of the future should be taken equally seriously. A localised present with a past defined by timelike worldlines cannot contain the information needed to predict the future. This is an essential part of the structure of spacetime.[16] This structure is integral to thinking about the causal structure of the world, among many other things. We should no more dismiss the limits of predictability due to the lightcone structure than we should dismiss its impact on causation. This is a limit on knowledge, but it is a limit set by the world and not by human capabilities.

Openness is indeed purely a feature of the presentist - internal - view on time. But it is an unfounded and undefended assumption that the external perspective should be taken as more fundamental than the internal perspective. It is in some ways a mere epistemic limit that we cannot ever access the facts about the future from the internal perspective and that we cannot predict them using present resources; but this limit is part of the intrinsic metaphysical structure of reality. To call this

---

[15] This has been largely the argument that the current literature on open futures has focused on. As we saw in the previous chapter much of this is based on Stein's (1968) relativistic notion of becoming based on causal connectability and the past lightcone. However, although recognised to be perspectival, it has not been explicitly mentioned that this is an *epistemic* limit as much as anything else.

[16] The branching spacetimes that Belnap (1992) and McCall (1994) developed as part of their open futures accounts have also been used to understand branching futures in the Everettian interpretation of quantum mechanics (see Wallace 2012 chapter 8 and Bacciagaluppi 2002). There is no genuine collapse in the Everett interpretation, which is why Mariani and Torrengo favour GRW; the future is perfectly predicted by the wave function (all futures are actualised). The branching of spacetime does not constitute the quantum probabilities. But the close connection here is undeniable and an account rooting the open future in quantum mechanics on these grounds could get off the ground. However, as I will argue in later sections we can get an open future without necessarily referring to quantum mechanics.

merely ignorance would be a mistake.

3.3 Perspectival Openness

Mariani and Torrengo give a clear statement of why one might reject epistemic openness:

> "clearly an epistemic take on the notion of unsettledness is uninteresting. If the future is open in a more robust and genuine sense, the unsettledness of future contingents have to be interpreted metaphysically, as an objective feature of the external world." (Mariani & Torrengo 2021a, p 3926)

The point of contention here is that epistemic matters are not objective and do not tell us about the external world itself, only our subjective capabilities. Alongside this is often found the assumption that all perspective dependent features are epistemic. Perspectival openness depends on our position in time and our practical abilities to make predictions. It is too subjective, or so the thought goes. I have introduced some introductory reasons not to be concerned about perspectival indeterminacy, but we can go further in responding to the charge that it is merely epistemic and subjective.

Two strategies for responding to this are available. The first, which has been the approach taken by Ismael (2016a; 2016b; 2023), is to embrace the charge of subjectivity by taking epistemic openness as a backdrop for explaining why we can control and influence the future. This accepts that the openness of the future is subjective but grounds it on the physical asymmetry of knowledge. In this case the fact that the openness is epistemic and subjective is unproblematic as the goal is to explain human experience and show how it falls out from the background physics governing the world.

But we can go further than this. I will argue that perspectival openness is an integral part of physics and that it is false to claim that it does not tell us anything about the objective, external world. This can act as a basis for metaphysical indeterminacy that goes far beyond just the indeterminateness of quantum mechanics laid out in the previous section. Perspectival structure is the basis for epistemic concerns, but we shouldn't be wary of it on these grounds. Agency accounts such as Ismael's are just special cases of perspectival openness and only work due to the metaphysical indeterminateness that is implicit in the background.

3.3.1 More than Epistemic: Agency Accounts

Ismael (2016a; 2016b; 2023) argues that the future is open due to paradoxes between how we represent the world and our role as agents in it (although this need not necessarily be construed as a *human* agent as opposed to an artificial one with similar capabilities). This works against the backdrop of the epistemic asymmetry. Our knowledge of the past (and lack of knowledge of the future) are the basis for the way we deliberate and act. However, the epistemic asymmetry on its own does not constitute an open future. She quotes Penrose (1979) as motivation for this conclusion:

> "the direction of psychological time...is not just a question of the past being (apparently) more certainly knowable than the future...It is not the ease in inferring the past that is relevant here, but the feeling that the past is unchangeable...it is not the difficulty that we might have in guessing...the future that concerns us, but the feeling that we can affect [it]..." (Penrose 1979 p. 594-596)(as quoted in Ismael 2023 p.7)

Her aim is explicitly to explain the *feeling* of temporal progression and why the openness of the future is an important part of that. Ismael is also a proponent of localised temporal becoming and this can be taken as a background to her views on the open future. Her focus, however, is not on the ontological side of temporal becoming but on how it relates to the experience of an embedded agent. The assumed background ontology seems to be the basic eternalist block without the reinterpretation of this that the previous chapter discussed. Although as we will see, Ismael argues strongly for the embedded, tensed, perspective but focuses on this in terms of an agent's experience rather than as an ontological claim.[17]

---

[17] The lack of explicit discussion of the ontological side of this picture should not be taken as an objection to the view I have developed. Even though all the mentions of ontology refer to the straightforward eternalist picture, Ismael says "In the case of time, we don't think of the future as there already, waiting to be experienced. We think of it rather as coming into existence as we experience it." (2016, p. 149) The idea of becoming as an ontological claim based on

An agent has to describe the world in order to theorise about it. And they inevitably have to include themselves in the description. Day-to-day physics, Ismael argues, avoids this as the goal is to model other systems. Even cosmology "maintains the imaginative fiction that we are sitting outside of the universe looking down." (Ismael 2023, p.1). However, a complete representation of the universe must include the agent embedded in it and from this comes paradoxes of self-reference. Ismael's example of such a paradox is to imagine a computer that can be asked any question about the physical world. It is equipped with the necessary knowledge of the state of the world and the laws of physics to be able to answer, and its answer will appear in its output channel. There is a simple question that the computer will be unable to answer truthfully regardless of what knowledge it has: 'Is the answer to this question that's about to be displayed in the output channel 'no'?'. The computer cannot answer this truthfully because its actions in the world (printing the answer in the output channel) interfere with its ability to make predictions. It cannot make these predictions without taking into account the way that it is representing and acting on the predictions. This leaves open a number of potential futures that the actions of agents in the world can bring about. Ismael labels the transformation of potentiality into actuality as becoming, i.e. the passage of time. This is, however, a limit on *representation* and not on the world itself; becoming occurs at the level of experience:

> "This isn't going to affect what the system can do at a *physical* level. It does place constraints on what the system can truthfully represent and it does lead to an essential incompleteness in the worldview of an embedded Agent." (Ismael 2023, p.9)

But this is still an important part of explaining how the world works and what our experiences of being part of the world are. Dismissing this as subjective and therefore uninteresting is neglecting a large part of what it means to explain the world. Additionally, Ismael shows that the incompleteness has a strong physical basis in statistical physics. She follows Albert's (2000) account of the arrow of time in statistical mechanics (and thermodynamics). Albert uses this account to explain the epistemic asymmetry (I will explore asymmetry and more details of Albert's

---

reinterpreting the eternalist block is well supported by this.

account in section 3.4 - for now I am simply concerned with the link between openness and physics in a general sense). The abundance of information about our macroscopic past is what makes it possible for agents like us to exist and to use information in the way that we do. Another consequence of statistical physics is that small changes made to the microscopic present propagate forwards and have effects in the future; but this cannot be done in reverse. This is dependent on the epistemic asymmetry; the abundance of macroscopic traces of past events constrain what the past could be and only large scale changes to the present that encompass all possible traces could lead to a different past being made possible. These large scale actions are outside the realm of any imagined agent's capabilities (precisely because records are so abundant and widespread). Although the openness of the future is part of subjective experience, the features that create this experience fall out of the relevant physics. It holds separately from any individual agent.[18]

Fernandes (2023), similarly working within the Albert framework, also considers the question of an agent's ability to control or influence the future but not the past. She takes this to be closely linked to Ismael's account but focuses on the statistical basis for it to give a more detailed account of how an agent processes information, deliberates on it, and then acts. This project builds on the ongoing literature from Albert and Loewer (for example Loewer 2007, 2023) on this topic that explores the statistical basis for causation, counterfactuals, and agency. Fernandes similarly concludes that the conditions that limit an agent to controlling/influencing only the future are a result of statistical mechanics.

All of this feeds into a picture where the future is open and subject to an agent's control. This follows directly from the epistemic asymmetry and other facts about statistical physics that determine how agents can act and what influence they can have. Although the openness is relative to an agent, it is a result of the relevant physics.

---

[18] Ismael's work, including her book How Physics Makes Us Free, gives a far more detailed defence of the physical basis for this view which I refer the reader to.

3.3.2 Epistemic is Enough: Predictability

Ismael's account, for all it relies on physics, puts openness at the level of representation and an agent's subjective experience. She dismisses the epistemic asymmetry between an open future and a fixed past as not enough *on its own*. Only with the addition of self-reference paradoxes and the possibility of an agent's action does she consider this as openness. The epistemic asymmetry found in physics plays merely a supporting role (albeit an essential one). It is certainly true that our experience of temporal passage goes beyond just a lack of knowledge of the future. The sense of control is hugely important and undoubtedly contributes to a complete and thorough explanation of our experience of time. But putting temporal passage entirely at this higher level is a mistake; the epistemic openness found in physics is hugely important in its own right. Ismael's narrow concept of agent-relative openness is just a special case of a much wider perspectival openness. The probabilities in statistical mechanics and the physical basis for information can be used to substantiate the metaphysical indeterminacy from section 3.2.

The accounts of metaphysical indeterminacy from section 3.2 largely amounted to claims about failures in our ability to predict the future from our position situated within time. This is an important sense in which information is used in physics. The laws together with the state of the world give us information about the future. Not all laws take this simple dynamical form but many do, and all our methods of testing theories empirically involve deriving dynamical effects and making verifiable predictions about what *will* happen under certain conditions.

The attempts to substantiate metaphysical indeterminacy rest largely on finding ineliminable metaphysical indeterminacy in the present on which the future depends (as Mariani and Torrengo do in their GRW account), or even looking at indeterminacy in which laws obtain (which Barnes and Cameron briefly consider).[19]

---

[19] This follows in the footsteps of Markosian (1995; 2013) who wrote on the problem of other times for presentism. The basic problem was how a presentist (in the original strong sense of the view and not the perspectival sense that I have been developing) can make any claims at all about times other than the present. These other times do not exist so how can they be talked about? While various solutions to this have been proposed, Markosian's proposal was that the present state is the initial conditions that can be used in conjunction to the laws of physics to specify events at all other times. I will not consider the other solutions to this problem, which include Bigelow (1996) and Crisp (2007) as notable examples, here as they are not relevant to

Assuming deterministic laws the present state along with the laws of physics fixes everything there is to know about other times and denies that there is an open future. Barnes and Cameron pick up on this and their strategy is to find ways around the determinism of laws and bring back uncertainty in predicting the future. As mentioned in section 3.2, special relativity places a huge constraint on this. A localised present does not have the resources to do this even in a fully deterministic world. An embedded, localised, perspective is necessarily limited in the information that can be used to feed into the laws of physics.

However, special relativity is not the only constraint on our ability to predict. These accounts consider only fundamental fully deterministic laws, which use precise and fully specified microstates to predict. Mariani and Torrengo respond to this by looking to the indeterminacy of quantum mechanics. Far more commonly, however, we make predictions using statistical mechanics or other high level theories, which also often make probabilistic claims about the future. On the same basis that quantum mechanical probabilities leave the future open, so too should these probabilities. It is these theories that constitute the epistemic backdrop to Ismael's self-reference paradoxes. These sorts of probabilities are not fundamental, which is why they are generally dismissed. But to do so is a mistake, and undermines their effectiveness in explaining the world. Similarly to special relativity we can view these sorts of probabilities as perspectival limits on the information available to make predictions. They are perspectival in the sense that they are indexed to which level of reality we are working from (rather than to the position).

Many accounts have been given to try and explain how we get thermodynamics out of statistical mechanics and what both theories tell us about the external world. Albert's (2000) discusses this and uses a past hypothesis and a statistical postulate to constrain the possible histories of the universe. This account explicitly introduces a probability postulate (the statistical postulate) that can be used to derive probabilities for macroscopic events at any times. It is from this that we get ordinary expectations such as the assumption that ice will melt and milk will disperse in coffee. Ismael uses Albert's account specifically as it relates to the epistemic asymmetry in the form of records and to the asymmetry of influence. However,

---

this project.

these are just part of the wider project of explaining the thermodynamic gradient. Wallace (2012) provides an alternative account that uses coarse-graining and the irrelevance of lower level details to justify the predictive success of macrodynamics. What is undeniable is that higher level theories are extremely predictively successful and have extensive empirical justification. And the predictions they give are probabilistic despite being derived from deterministic laws.

There is significant debate around how to interpret the probabilities both in these accounts and in statistical mechanics more generally. One claim is that the probability distributions represent our inability to determine the exact microstate, making them entirely epistemic. Without this knowledge we are forced to use probability distributions that represent our uncertainty in the exact state. However, several accounts of probabilities have been developed that aim to show that probabilities are objective parts of the world, even if they have an epistemic gloss (often calling them epistemic chances or objectified credences). Ismael (2009) argues that probabilities are relative to well-defined physical macrostates and that they are an indispensable part of physics. Myrvold (2021a) and Strevens (2011) base their accounts of probabilities on features like the dynamics or ineliminable noise at the microscopic level that makes the exact initial microstate - or even the initial probability distribution we assign to a system based on our knowledge - irrelevant to predictions. All these accounts show that probabilities are objective and latch onto real physical features even in a world governed by deterministic physics at the fundamental level.

It is also possible that probabilities in statistical mechanics come from quantum mechanical probabilities. This is an open question (see Popescu et al. (2005) and the accounts mentioned above for more details). But if true, harkening back to Mariani and Torrengo's account in section 3.2, this would lead to an open future essentially based on quantum mechanical indeterminacy in the present (but with a much wider understanding of probabilities and how they feed into higher level theories). There is still debate about how to interpret quantum probabilities but they are generally taken as the most objective probabilities one could get in the world.

To relate this back to openness more explicitly, statistical mechanics cannot specify a single unique future relative to an initial macro condition but instead gives a range of possible futures, each with a different probability. In these cases we can use the accounts of metaphysical indeterminacy from section 3.2: the truth values of future claims are not determinate. The potential challenge in doing this is the non-fundamentality of the probabilities in statistical mechanics. The accounts of metaphysical indeterminacy are not necessarily straightforward to apply to questions about levels in science (this is likely part of why Mariani and Torrengo jump to utilising quantum mechanical probabilities rather than any others). In the language of truthmakers it seems if there is anything at any level that makes a claim true then the truth value of that claim is fixed. For Barnes and Cameron if the future is predictable by the fundamental laws combined with the initial state then it is settled. This is a reductionist approach. Emergent laws and probabilities, however, could be relevant when we have a localised present perspective and there is not enough information to accurately use the fundamental laws. Emergent laws are often particularly applicable to semi-isolated subsystems.[20] Wilson's account provides better resources for dealing with levels. A determinate for a determinable is a state of affairs that obtains in the world. A state of affairs is a very general notion and we can easily conceive of the state of affairs being described at different levels. This would allow a claim to have a determinate at one level but not at another. Wilson explicitly discusses this sort of level dependency.[21] For example we might have a determinate that a certain object is red but not a more fine grained determinate as to whether it is scarlet or crimson. Similarly a macrostate

---

[20] Only considering subsystems does not *rule out* using fundamental laws as these too can be applied to isolated localised systems, emergent laws will also be constrained to the information in the localised present. However, emergent laws have better apparatus to deal with localisation. Ismael (2009) discusses the thermodynamics of open subsystems in her account of probabilities. When fundamental laws are used there are normally a number of assumptions and idealisations (such as negligible noise, no air resistance etc) that have to be made to define an effectively isolated system and treat any outside influences as negligible. For any complex system even semi-isolation is unlikely and emergent laws are specifically designed to recognise that the system is open but disregard irrelevant lower level details and minor influences from outside the system (see Franklin & Robertson 2022; Wallace 2012). Along with thermodynamics another example of this is decoherence, which is explicitly an open systems theory. Chapter 5 will discuss localisation in quantum decoherence.

[21] Mariani (2022) also discusses this and more generally the possibility for non-fundamental indeterminacy, although no physical examples are given.

in thermodynamics plays the role of a determinate for claims about the future, but it is indeterminate as to precisely which of the possible futures compatible with the macrostate will be the actual one. More likely than not the microstate of a localised present will not contain enough information to act as a more detailed, finer-grained, determinate than the macrostate to resolve this indeterminacy at a lower level (for example the microstate may be too sensitive to influences from outside of the localised present). And even if it did, this should not undermine the objectivity of the higher level probabilistic indeterminacy based on the macrostate.[22] Whether the future is open - whether it has determinate truth values - is relative to the set of information in the present. This depends both on the level of description as well as position within time and the localisation of the present.

One might still be tempted to dismiss this as only pertaining to higher level theories and not revealing genuine objective metaphysical features of the world. However, one should only be willing to do this if one is also willing to dismiss statistical mechanics and thermodynamics (as well as the multitude of other higher level theories that use similar methods).[23] Certainly some may be willing to do so (thermodynamics has often been accused of having an epistemic basis) but the objectivity of thermodynamics and the value of other higher level theories in providing novel explanatory value that cannot be captured by the lower level has been extensively defended in the literature on emergence and reduction in physics. This is at odds with viewing thermodynamics, and its probabilities, as merely epistemic. Of particular interest here is the defence offered by Robertson and Prunkl (2023) that picks up on the information theoretic basis to thermodynamics. They argue that a realist and objective interpretation is perfectly possible because, even in places where it explicitly uses epistemic concepts, these are taken to link onto objective physical structure in the world. They conclude that:

---

[22] I would also argue that, even in a case where we did have a microstate in the localised present that was perfectly isolated and contained all needed information to predict the future with deterministic micro laws, we should still maintain that the future is indeterminate at the higher level. This would be the reverse of Wilson's situation as we would have a determinate at the fine-grained level but only probabilities at the coarse-grained level.

[23] Even in quantum mechanics interpretations such as the Everett interpretation rely on emergent structure to interpret probabilities. It is also increasingly common to question what it means for a theory to be fundamental and if there is a fundamental level of reality at all (e.g see McKenzie 2011).

> “Thermodynamics might not be fundamental, but it is objective.” (Robertson & Prunkl 2023, pg 9).

In just the same way we can reject Mariani and Torrengo’s accusation, quoted at the beginning of section 3.3, that epistemic openness, even with a metaphysical gloss, does not tell us anything about objective features of the external world. The epistemic asymmetry and the openness it describes *is* objective and constitutes real physical structure in the world. It is perspectival and level dependent; it does not necessarily feature in the fundamental laws of physics. But this should not be reason to dismiss it any more than we would dismiss thermodynamics. At the level of statistical physics the future is genuinely and metaphysically open.

A standard eternalist (who has only the tenseless block of facts) can explain probabilities only by viewing them as resulting from restricting the total amount of information available in the block universe (i.e. we restrict to the subset of facts available at that spacetime point at the relevant level and predict from there). This does not necessarily mean they cannot view emergent structure as objective and important - the eternalist can certainly accommodate indeterministic laws. But it makes the accounts of probabilities much weaker and more epistemically inclined. Recognising the future described by these probabilities as genuinely metaphysically indeterminate brings forward the very real and influential structure that they capture. The result of this structure is that future statements have genuinely indeterminate truth values from certain perspectives. Acknowledging perspectives as objective and unavoidable aspects of reality makes it much clearer why emergent theories have so much novel explanatory value. The statement ‘x occurs with probability 1/2 at t_2’ is made true by the current macrostate that defines this probability. A standard eternalist would have to explain how this can be true despite ‘x does occur at t_2’ also being true (made true by the actual state of affairs at t_2).

3.3.3 Epistemic is Enough: Physical Information

We can go deeper into this attempt to justify the objectivity of epistemic structure, picking up on the claim by Robertson and Prunkl that the epistemic basis of thermodynamics captures objective structure in the world. Predictions and probabilities are only part of a much wider sense that physics uses information, even if these are the most obviously related to the openness of the future. Information is also directly encoded in physical states or analysed through the process of measurement. Recognising information as physically realised goes a long way to understanding why probabilities and predictions based on levels and perspectives are objective. It also shows why picking a perspective is not the same as arbitrarily restricting the total information that is available in the block universe; instead, perspectives are naturally defined based on physical structure that imposes these limits in specified and lawlike ways.

Encoding information in physical states is an important, but largely overlooked, aspect of Ismael's self-references paradoxes. The archetype paradox that she uses is a computer printing the answers to questions. It is the way the output screen - in itself simply an arrangement of coloured pixels - represents information that makes it a paradox and prohibits actions in the world. The computer, if it is bound to be truthful, physically cannot print an answer to the question 'Is the answer to this question that's about to be displayed in the output channel 'no'?'. This is more than a paradox in information but is a constraint on what correlations can be formed in the world (the correlation between output and the processing the computer does to produce the correct answer). In computability theory, how to physically implement these sorts of algorithms on a physical dynamical system is an active field of research. (For example, see Prokopenko et al. (2019) that looks at undecidable dynamics versus chaotic dynamics in relation to computational self-reference paradoxes. Also Hernández-Orozco et al (2018).)[24] Ismael admits that her openness is at the level of representation but overlooks how strongly representation is a physical process.[25] This picks up on technical concepts of information that information theory (both quantum and classical - as well as large parts of computer science) are built on. Timpson (2013) analyses this concept of information. The basic idea behind the formal theory of information is Shannon information, which

---

[24] This field is part of a wide exploration both of the philosophy of maths and the physical basis of life. Interestingly there is also a strong connection to emergence discussed in relation to these sorts of dynamics, where emergence is defined in this field as underivability of the macroscopic future state from the initial microscopic conditions. This is very similar to the probabilities of thermodynamics and is also closely related to probabilities in chaos theory.

[25] Of course Ismael's paradoxes do not necessarily involve printing information, and actions that interfere with predictions encompass a spectrum of physical processes. This is only an archetypal case.

is a way of characterising the possible states a system can have that can be used to represent information. For example a system with 5 different possible states can be used to represent 5 bits of information (each state representing a certain bit of information). This can be built on to characterise correlations between systems. Information theory uses these basic concepts to deal with complex processing tasks. Timpson is careful to distinguish this from the everyday sense of knowledge. The technical notion of information is a way of using physical systems to represent and perform tasks with information. It does not describe (directly at least) how we learn about the world or what knowledge we might get from seeing a particular message. For example, information theory can describe the physical requirements needed to send a message of a certain length and complexity (for example an 8 bit code sent in binary or ASCII). However, the knowledge a person might gain from the message encoded will depend on their ability to understand the code, the content of the message in ordinary language, as well as the context of the message and how its fits into the variety of background knowledge the person might have about the subject being communicated.

When thinking generally about epistemology and our ability to know things as agents we tend to mean the everyday sense of knowledge. The two concepts are connected but distinct. Technical information describes how states and correlations between systems work in the physical world. This does not give a quantitative description of the everyday information we might have, but it does constrain the ways that it is possible to manipulate the world to get or use information. It is these features of information - and not knowledge - that Robertson and Prunkl (2023) are relying on to justify the objectivity of thermodynamics. And it can equally be applied to objectify the limits that are placed on the predictability of the future. There are certain states of the world that can form certain kinds of correlations and these are the basis for extracting, representing or using information. These states can act as determinates for predictive claims in Wilson's account of the indeterminateness of the future.

Not all epistemic ignorance indicates a metaphysically open future. The future is not open because I have not yet glanced at my measurement results or because I have not understood the significance of something I am seeing. But the future is open when the world itself places limits on what information is encoded in a particular perspective. When it comes to this sense of information there is no meaningful difference between metaphysical indeterminateness and epistemic indeterminateness - they both have the same physical basis. Predictability is a physical constraint on what the dynamics of the world are and what sort of correlations can form as much as it is an epistemic constraint.[26] And methods of prediction, observation, and control take advantage of specific features of our theories (and the external world they describe) and often reveal important structure in them. This has nothing to do with an agent any more than 'observers' being used as a synonym for reference frames in special relativity means that relativistic effects are subjective.

3.4 Asymmetry

If the future is open, then why is the past fixed? This question has received comparatively little attention, yet to truly understand how openness plays a role both in physics and in the flow of time it is unavoidable. It is generally expected that while claims about the future can be indeterminate, claims about the past are fixed. Without this the future is no different from the past; there is no sense of a transformation from one to the other that gives time its distinctive dynamic appearance (see chapter 1). Likewise within physics, without getting a handle on what is *not* probabilistically indeterminate we cannot properly work out where and why objective probabilities do arise.

Both localised temporal becoming and the open future model of time considered in the previous chapter assume the determinateness of the past. Passage is the coming into being of the present moment and as it comes into being it makes definite the previously indeterminate future. At the next moment a new present comes into being and, to avoid a collapse to the growing block view, the previous moment must cease to exist. In this sense becoming is perfectly symmetric; and following from this the indeterminateness of the future should be mirrored with an

---

[26] Ismael's and Fernandes' justification of control over the future depending on epistemic openness from statistical physics is along these lines but does not go nearly far enough in recognising the physical basis for abilities as agents. And they also do not appreciate how this goes well beyond agents and is a much wider constraint on physical systems of all sorts.

indeterminate past.

Views on asymmetry differ between the different advocates of temporal becoming. Dorato (2006, 2000) rejects symmetry and assigns an intrinsic arrow to the process of becoming based on the principles of causal connectibility in Minkowski spacetime. (Dorato rejects the open future thesis and so has no questions about what fixes the past. His arrow is merely to give directionality to the processes of physics). He claims that:

> "Temporal becoming involves the issue of an asymmetric causal or temporal relation, and therefore the complex question of the direction (arrow) of time, but has nothing to do with allegedly ontic asymmetries between past, present, and future" (Dorato 2006, pg. 572)

However, it is unclear how the metaphysical process of becoming inherits its arrow from spacetime structure. Are these two separate arrows that simply happen to align? If the connection is deeper than this then how exactly does becoming relate to this causal/temporal relation? These questions are left unexplored. The open future model similarly relies on the structure of spacetime and the assumption of asymmetry between the past and future lightcones. While the lightcone clearly limits predictability, and as such is a suitable source of openness, it does not give any *physical* structure to fix the past. Others such as Rovelli (2019), claim that becoming itself is symmetric and any asymmetry is found purely within physics. This leaves becoming a much weaker concept that seems to play little role in explaining reality.

The discussion in this chapter suggests another answer. The openness of the future is a result of the structures in physics that place limits on knowledge of the future. Likewise the fixity of the past must depend on the structures that ensure knowledge of the past. Hence the past is determinate because of and conditional on the records we have of it. Records will tell us about what is in the causal past, but are a physical realisation of this information. This is largely in line with Rovelli's claim of fundamental symmetry with asymmetry derived from physics. However, combined with the discussion from the previous sections we can see that this goes beyond just directedness of our physical laws; is a derived *metaphysical* asymmetry.

The change in reality from moment to moment that becoming describes is, in itself, symmetric. But the structure of reality is such that the content of what becomes at each moment produces an asymmetric structure.

In the following two sections I will first argue that fixing the recorded past satisfies all of our everyday intuitions about the past (including the impossibility of backwards causation) and that this is integral to the asymmetry found in physics. This ties the asymmetry of flow directly into the reductionist project for explaining time's arrow. And second I will lay out what the resulting metaphysical picture would look like when using records in conjunction with localised temporal becoming as an account of passage.

3.4.1 The Recorded Past is Fixed

There is a distinct asymmetry in our knowledge of the past versus the future. Our knowledge of the future tends to be probabilistic and uncertain. On the contrary, we are easily able to make inferences and claims about a past that we consider to be fixed and definite. Prediction and retrodiction are, in themselves, fairly symmetric; the asymmetry is primarily due to *records*.[27] These are extensive traces of the past such as photographs, fossils, footprints etc - anything that bears a direct correlation to a past event and allows us to make inferences about what has happened. Records are a prominent feature of our experience of temporal asymmetry. We can also draw on the way that records are connected to many other temporal arrows such as causation and counterfactuals to understand why the fixity of the past should depend on records.

The research program spearheaded by David Albert and Barry Loewer for explaining thermodynamic asymmetry, and its connection to causation, agency and other arrows, relies on records as a central part of the argument. They first appear in discussions of Boltzmann brains. Our evidence that the present universe has *not* merely fluctuated into its current state from a previously higher entropy state is the assumption that our records of a low entropy past are veridical. This provides justification for postulating the past hypothesis, which is needed to explain these

---

[27] The next chapter will discuss this, and what exactly a record is, in far more depth, as well as looking at questions about how records are localised, macroscopic systems.

observations (see Albert 2000 for details of this program). In turn, this then allows us to explain how asymmetric records arise.

This can be more specifically linked to the probabilities discussed in section 3.3 (several of the probability accounts explicitly align with Albert and Loewer's program). The probabilities are of future evolution given the present macrostate. Records, along with the past hypothesis, significantly limit the possible past macrohistory of the universe. Although they don't precisely specify the past microstates, they give a much narrower probability distribution over the past macrostates than they do over potential future macrostates that the system could evolve into. The set of lower entropy histories is much smaller than the set of higher entropy futures - i.e. the past is fixed in comparison to the future.

Beyond this, records then also play a crucial part in explaining the counterfactual and causal asymmetries, which give us our intuitive ideas about what the past being fixed means. The idea of using records for this predates Albert and Loewer's program and stems from Lewis (1979)'s analysis of counterfactuals. The basic idea is that records block backwards counterfactuals because the event the record is a record *of* must exist in the past. Records generally give highly specific information about the events they are produced by. Generic events do not have this property, they can often be caused by multiple different causes in different circumstances and the connection may be hard to discern. Records, however, are states that carry specific details of the event and can be linked back to what created them much more clearly. If we alter the event in the antecedent to test a backwards counterfactual but keep the rest of the present the same then the abundant records that exist in the present ensure that the past (the consequent) remains as it was. So changes to the antecedent do not alter the consequent, meaning the counterfactual does not hold.[28] Only by altering *all* the records that exist of a past event could the past be changed. While this was first applied to counterfactuals, the basic idea can be differentiated to support a number of subtly distinct arrows including the more general causal arrow and the arrows of deliberation, influence, and control. Jointly the arrows of deliberation, control, and influence - which closely match counterfactuals and

---

[28] In Lewis's framework this essentially means that in the closest possible world in which the truth value of the antecedent is flipped the consequent does not change.

causation but narrow down on the ability of an agent to knowingly and deliberately affect the past - ensure that the past is fixed as we expect it to be (see Fernandes (2023) for detailed discussion of this). This allows some leeway for minor backwards influences in cases where, for example, no records at all exist. But we can never knowingly manipulate the past to our own ends. In terms of our everyday intuitions this is largely why we consider the past to be fixed in a way that the future is not. Nothing we do can alter it, and knowledge of it is presented to us as a done deal. The fixity of the past is not an absolute and inviolable condition however, it simply appears that way to the average person because records are so abundant that exceptions never arise.

All of these arguments largely belong to a single interconnected research program (albeit an expansive and complex one). But records also feature more generally in other possible approaches to understanding causation. For example Dummett's bilking arguments about causation rely on whether or not an agent has relevant knowledge. The Bilking Argument is a paradoxical set up in which an event C causes an earlier event E. However, an agent could interfere with the process *after* E has taken placed and prevent C from occurring. This leads to a situation in which the C cannot have caused E because E occurred without C. The agent has 'bilked' E of its cause. Dummett (1954; 1964) identified that the problem in this situation comes about because we have knowledge of the effect and can therefore act to interfere. It is records and information that make backwards causation problematic here.

Price (1996) considers Dummett's solution to the Bilking Argument: that backwards causation is only coherent if we do not have knowledge of the effect before the cause occurs. In a case where we could not be sure that E had occurred until after C had also occurred there would be no potential problem. This conclusion leads Price to argue that in order to coherently evaluate counterfactuals about backwards causation we must hold the *accessible* or *knowable* past as fixed (but not the past as a whole).

These arguments, although involving what an *agent* can do to the past, are strongly based in physics and are independent of any particular agent. The way that records are correlated to the past limits what effects actions in the present can have. Where backwards causation is acknowledged to be possible, it is constrained to affecting only the events in the past for which there are no records.

Furthermore, even beyond the physics used in these accounts to explain what control and influence really amount to, many other areas of physics should lead us to doubt that the past is fixed beyond what is recorded. Another area that uses ideas about records and inference in explaining asymmetry and the fixity of the past is quantum mechanics. While I do not want to base openness entirely on quantum mechanics as Mariani and Torrengo do, it is certainly not insignificant to consider how quantum openness fits into this picture (later chapters will do this in more depth). It was recognised early on that the limits that quantum mechanics places on predictability apply to the past as well as the future:

> "the principles of quantum mechanics actually involve an uncertainty in the description of past events which is analogous to the uncertainty in the prediction of future events...this uncertainty in the description of the past arises from a limitation of the knowledge that can be obtained by measurement of momentum." (Einstein, Tolman, & Podolsky 1931, pg. 780)

The uncertainty principle places limits on what information can be accurately determined and as such we are unable to completely retrodict the past based on present information. Di Biagio, Dona and Rovelli (2021) present an argument following this that the time asymmetry of quantum mechanics as a whole is entirely due to the asymmetry brought about by the process of making inferences, while the underlying dynamics are symmetrical. While there are many other possible explanations available of quantum asymmetry, it is undeniable that measurement and inference are important factors.

This is not an exhaustive list of everywhere where records are used but it indicates a number of important places where records can be linked to the claim that the past is fixed in a way that the future is not. Certainly it shows that records are what give us our commonplace assumptions about the past and suggests that where records fail we should give up our assumptions in favour of the results that physics has provided. The following three chapters will develop the physical basis for this further so I will let the argument rest for now.

3.4.2 Presentism, Flow, and Records

I will now consider how this translates into a metaphysical difference between past and future.

Section 3.3 discussed how states of affairs, such as macrostates in thermodynamics, can act as determinates for claims about the future based on how they define probabilities for future events. Similarly when a state contains records of the past these act as truthmakers for claims about the past. Because present states (which are localised and can be either macrostates or microstates - although localised macrostates are far more likely to contain the information needed to make effective claims about the past and future) commonly contain a large number of records but only define probabilities of possible futures we get an asymmetry between past and future. Future claims are largely indeterminate while we have many determinate claims about the past. This is possible precisely because there are so many abundant records of the past and because they contain fairly precise information about the past. Additionally, where a single record is not informative enough to precisely fix the past, a set of many different records formed by the same past event can support and verify the information contained in other records.

Using records to build on presentism and explain how we can ground truths about the past and future on only presently existing things has been proposed before: Kierland (2013) makes a similar suggestion of solving the grounding problem for presentism using ‘records’. Based on the intuition that “reality carries its history along with it” (p. 180) he proposes that there should be ‘records’ of the world’s history existing in the present world that satisfy four constraints: (i) the record should be an effect of the past, (ii) the record must be perfect (i.e. carry complete information about the history of the world), (iii) the information in the record must not require significant mathematical/logical work to recover, and (iv) the record’s sole purpose should be to be a record of the past. Kierland dismisses the idea of using features of the present world (such as the physical records I have talked about) as this sort of record is only able to satisfy the first constraint (arguably

(iii) as well but often records are far from straightforward). Instead he proposes introducing specially made 'records' that are metaphysical properties of the present world that are dedicated to exactly this purpose and no other. This is along the lines of common presentist solutions for how to talk about other times. The most common solution is Lucretianism (Bigelow 1996; Crisp 2007) where the present contains past and future oriented properties that ground past and future oriented facts. Kierland's records are a variation of this. The obvious ontological downside of this is, of course, having to posit a new kind of property in the world. The world is now no longer made up of presently existing things but also past and future oriented *properties*. The relation of these properties to existence and the actual present state of the world is unclear.

Several of these constraints are unmotivated, however, when we consider the perspectival ontology that I am using; in this context actual physical records do the job perfectly well. A standard presentist has no other resources than the present and the whole concept of temporal extension with a past and a future as well as order, flow and a fixed history must be captured within this. Because of this the records must be complete (perfect), otherwise time would have unexplainable gaps in it. Constraints (iii) and (iv) come in because the record must be a dedicated way to understand what other (non-present) times are and as such they cannot be reducible to something serving another purpose in the present world. Without this there would be no distinct concept of *the past* as something separate from the present. For a non-standard presentist the record does not have to do all of this. Other times are already taken to be as real as the present – just only from their own perspective or from the external perspective – so we do not need to create the concept from scratch. As for constraint (ii): the tenseless perspective is a way to capture the complete and fixed history of the world so this no longer needs to be entirely contained within the present record. Records are abundant enough to fix the past relative to the present to the extent that we would expect. In the rare cases where there are no records, or the records only specify the past event to a certain probabilistic degree, then we should accept that these aspects of the past are indeterminate in the same way the future is. They still exist at their own locations in time and still have bivalent truth values, but are indeterminate with respect to the present. This is in line with the accounts of causation and control laid out in the previous section that show why the small possibility of backwards causation should not worry us. In this we should recognise that our intuitions that the past is universally fixed should be weakened to match physical accounts of control.

It is only the first constraint that has any motivation. We do indeed require that records are caused by the past. Even in this we can put aside a strict notion of causation requiring a metaphysical account of what this means, and simply say that records must be genuinely correlated to past events. This is important to rule out spurious or misleading 'records' that look like systems which provide information but are actually false. For Kierland this is done by definition of what his record properties are, but interpreting physical systems in the actual world is more complex. This, however, is once again easily taken care of by the eternalist perspective. We do not, in the present moment, need to be able to perfectly distinguish true from false records. Records are only truly records if they are correctly correlated to the past; this correlation is revealed in the eternalist perspective where cross-temporal relations such as this can be seen. Records are a certain type of physical correlation found in the world independently of our ability to correctly read and interpret the information. And while the correlation to the past event is cross-temporal, the existence of the record itself can be found in just the present moments.

Records within the present act as grounds for claims about the past to the extent to which they specify information about those past events. And most importantly we can ground these facts without having to refer to anything except presently existing systems. I have already discussed how this may have unintuitive consequences when it comes to areas that we have no records of, and shown that this is in line with how we understand the physical basis for backwards control and causation. Likewise it is worth noting that the future may be *more* fixed than expected. Although we have a distinct epistemic asymmetry between past and future, it is not the case that we know nothing at all about the future. To the extent to which our laws and current facts allow us to deterministically predict the future, we should consider those aspects of the future to be as fixed as the past is. The account I propose, however, provides a strong physical basis on which to base the fixity of the past and the indeterminateness of the future where it does exist. What is more, it naturally ties into the way in which many temporal arrows are explained through physics. One failing of many dynamic views of time is that they postulate fundamental metaphysical flow and directionality but fail to show how this produces the everyday experience of time that we have. Our everyday experience rests on temporal arrows such as control and our experience of causal relations as well as the ever present epistemic asymmetry. Building these arrows directly into the metaphysical picture of flow answers this question.[29]

Another somewhat unintuitive result is that indeterminateness and fixity are level-dependent. This connection will becoming increasingly clear in the chapters to come; but it is already clear that openness and fixity primarily concern the macroscopic, emergent level of reality described by theories such as thermodynamics. The next chapter will also explain why records are macroscopic (a conclusion that is already assumed in the way that records are used in the literature on statistical mechanics and time asymmetry laid out above). Localised temporal becoming as a metaphysical model does not presuppose level dependence; becoming occurs between microscopic events just as much as it does between macroscopic ones. However, what makes becoming a substantial account of *flow* is how it can be connected to the the metaphysical asymmetry of openness and fixity. It is not just becoming at each moment but the *becoming actual* of an *indeterminate future.* The fact that openness is a result of level dependent epistemic structure means that the flow of time is likewise level dependent. At the microscopic level we have becoming but it is a fairly trivial notion and the presentist perspective does not have much value. However, at macroscopic levels, which are rich with epistemic structure, the openness of the future plays an essential role and the flow of time is an ineliminable feature of reality.

Of course some sources of openness - such as the lightcone structure - apply at all levels. Part of the argument of this chapter is that openness does not have a single

---

[29] There may be additional stages to recreating our full psychological experience. I do not, for example, claim to have explained what it means for us to *feel* dynamic movement through time (i.e. what our qualia for this is). Merely to have put in place, with a clear physical basis, all the features needed so that this is a small extra step rather than creating an psychological illusion of flow from nothing but the bare facts of the eternalist picture.

source but there are many cases which constitute failures of predictability. Openness comes in degrees and varies widely across different perspectives. Therefore there is no single unique flow. As special relativity has already forced us to accept the localisation of flow, this is a natural extension of this.

Adopting a level dependent model of time's flow takes seriously the extensive work on emergence and specifically the emergence of time asymmetry in physics. If asymmetry is emergent then flow should be as well; these two claims go hand in hand. Indeterminateness should not be dismissed just because it is not necessarily fundamental. And the world contains far more indeterminateness than we often recognise. The piecemeal spread of information across time and across levels reveals how important the processes of measurement, prediction and retrodiction are to understanding the world. This rich structure provides a map between different parts of reality and shows us how different systems and different times are related across different levels.

3.5 Conclusion

While the previous chapter laid out the ontological foundations of localised temporal becoming, this chapter has shown how this can be extended to give a complex reality structured by the indeterminateness of the future and the fixity of the past. Leaving this out of our complete picture of reality means ignoring many important and salient parts of physics that we use to explain a large array of physical features. It is easy to dismiss indeterminateness as epistemic but doing so means ignoring the structure of the world that forms the foundation for how we can learn about past and future times (as well as distant events). In this sense the view of the open future that I have presented here has a far more inclusive physical basis than the views of the open future most commonly considered. The open future literature is hugely focused on propositions and semantics in everyday language; where a physical basis is considered it is largely constrained to special relativity or to fundamental quantum probabilities. However, the view presented here provides a new way to think about how the open future is found in many of our physical theories. The lightcone structure is just one example of an epistemic limit.

Indeterminateness, and fixity, are features of time that are relative to our perspective inside of time. They strongly relate to our subjective experiences, but the basis for this is objective physical features of the world. The process of becoming describes how these physical features change over time. As each successive moment comes into existence it contains new records of what came before - keeping things fixed even when they cease to exist - and closes off potential futures that are no longer compatible with the newly created physical state. It is the physical state of the localised present that grounds claims (either definite or probabilistic) about the past and future.

Flow, and openness, is both perspective and level dependent. This is a natural result of understanding flow to be intimately connected to the way that different moments are linked together. Different theories make predictions (and retrodictions) in different ways, using a variety of dynamical equations as well as other techniques; correspondingly there are many sources of openness, such as the light-cone structure or probability distributions. Pinning down exactly what claims are fixed and which are open in any given perspective is no easy task and I do not take the work here to provide a definitive answer for how to do this. I have described a general strategy of grounding this on the information available to make inferences and provided strong motivation for why theories of metaphysical indeterminacy *need* to take this into account, without giving all the details as to how. A full theory of semantic truthmakers, a more precise look at the different accounts of metaphysical indeterminacy, and how this relates to states of affairs in the world (which Wilson's determinable account uses) must be adapted to deal with level dependency and perspectivalism; but this is beyond the scope of this thesis (I have attempted to indicate through footnotes in this chapter some points where this might touch on what is laid out here and provide references to literature which focuses on these sorts of questions). The resulting theory would have to account for a complex range of factors; but this is an unavoidable result of acknowledging the full extent to which the propositional claims we make are theory laden and make reference to level dependent ontologies.

Of course it is essential to this account that we have a thorough understanding of exactly what the epistemic structures in the world are. This chapter has laid out various areas of literature that help elucidate this. The remaining chapters of this thesis will continue to examine how physics uses records in particular so as to better understand how they can act as a basis for our metaphysics.

4 A Non-Idealised Account of Records

4.1 Introduction

Records of the past are a prevalent and inescapable part of our world. But there is no widespread consensus on what records are. The previous chapters have established how we can use epistemic structure in physics - of which records are a key example - to explain the metaphysical asymmetry between past and future. As such, a clear definition of what a record is, what features they have, and how they relate to prediction and retrodiction is needed. The aim of this chapter is to provide answers to these questions, as well as to show the importance of records in statistical mechanics and the physics of time asymmetry. The current most successful account of records is given by Albert (2000; 2015) as part of his wider analysis of time asymmetry. However, his account is limited by overidealisation that, while justified within the overall project he is aiming for, prevents using it more generally to explain the use of records in our physical theories. A non-idealised account is needed if we are to understand records and explain how they play a role in our theories.[1] The lack of such an account has already caused confusion in the philosophy of physics literature (more on this later).

I will give a new account of records that focuses on their practical use and explains how we make day-to-day inferences about the past. The pitfalls of overidealisation are well known and I will focus on what a non-idealised account can give us by centering the requirement that records must be *robustly* correlated with the past state of some other system. This defines what a record is and distinguishes it from retrodiciton; with retrodiction there is no unique correlation between events in the past and the current state, so any information about other times requires a calculation of the evolution of the present state either forwards or backwards under dynamical laws. For records there is such a correlation, which shows a direct depen-

---

[1] Of course no account is completely without idealisation. And in this chapter I will discuss many methods such as coarse-graining and statistical methods which are in themselves often considered forms of idealisation. But this account is non-idealised in comparison to Albert's and recognises that we cannot just pretend we are in an idealised situation where influences such as noise do not exist at all. Instead we must find rigorous ways of accounting for these effects. The setting, if not the methods, is non-idealised.

dency between the current state of the record and a specific past event regardless of the intervening evolution. And this correlation is robust against the influence of noise; this is what makes it identifiable and informative. Incorporating noise into the definition of records, and not just as an afterthought concerning practical implementation, will not only give a more practical account but will also highlight important features of records that have gone largely unrecognised: namely that records involve the use of *redundancy* and that records should be treated macroscopically when we theorise about them (by which I mean they must be modelled using the array of techniques such as coarse-graining, probability distributions etc that we commonly use to model macroscopic physics).

The conclusion that records should be treated macroscopically is valuable for a number of reasons. First, it explains the way that records are used in the literature on time asymmetry. So far, most of the attention given to records in the literature has followed Albert in considering primarily their role in explaining emergent temporal asymmetries. The knowledge asymmetry underpins much of our conceptual understanding of the past and the future and often plays a role in understanding agency and other directed human phenomena (for example Ismael 2006; Loewer 2023). Records and the associated asymmetry also feature in explanations of many other areas. Albert (2000) and Wallace (2012) use records to justify postulating their different versions of a past hypothesis. Loewer (summarised in 2023) grounds his account of counterfactuals and causation on the asymmetry of records. One important assumption is made throughout this literature: that records are macroscopic. This assumption is lightly made and does not amount to much more than saying that records are in the realm of everyday, observable things[2]; but it is important to a lot of their accounts (for example in Loewer (2023) the existence of records in macrostates but not microstates is essential to establishing free will) and Albert's account of records provides no explanation of it. Explaining why records are macroscopic and what this means for the way we treat them theoretically is important for these projects and for fitting records into the general picture of macroscopic asymmetry and microscopic symmetry. My account of records can do

---

[2] Huggett (2023) continues this assumption. Additionally, Hemmo and Shenker (2012) argue that measurement, upon which records are based, is macroscopic. Stradis (2021) argues that records must be observable.

this. It also suggests additional avenues for exploring the record asymmetry such as a connection to the irreversibility of decoherence and coarse-graining methods used to derive emergent asymmetric macro-dynamics.

Secondly, connected to but going beyond questions of asymmetry, memory systems are used in other areas of thermodynamics and questions about their scale and reliability have proved to be pivotal. For example, in thermodynamics there is the long standing debate about the thermodynamic cost of the measurement and erasure of information through Landauer's principle; Szilard's one molecule gas memory device is central to this literature. In nano and quantum computing preserving information through memory devices is an important practical challenge to computing at the microscale. I will look at how the literature surrounding Szilard's one molecule gas memory device points to how a lack of understanding of records can lead to their misuse and contribute to theoretical confusion.

4.2 Records as Measurement: The Ready State Account

Many theories to explain records have been proposed. Of these Albert's (2000; 2015) is by far the most detailed and successful and it is the only one I will explore in depth.[3] This section will lay it out and consider how overidealisation limits the account. Albert identifies records through the mechanism by which they give information; a process he calls inference by measurement. According to him there is a sharp distinction between this and retrodiction - the process of gaining information about the past by taking the current state of the system and evolving it backwards.

All systems are informative in a trivial sense that they can tell us about their own present state. This is the kind of information that is inputted into retrodiction. Records, however, additionally encode information about a past state they are correlated to. This sense of encoding information is familiar from work on information theory and physical information. The view in this field is that all information

---

[3] The other well-known account of records comes from Reichenbach who defines records as subsystems with low entropy compared to their surroundings. Earman (1974) explores objections to this account. Reichenbach's account is a backdrop for Albert's understanding. Other accounts include computational accounts (Hartle 2004; Hawking 1994; Schulman 2005), and the fork asymmetry account (Horwich 1988; Stradis 2021). Wolpert (1992) also produces an account that has many similarities to Albert's.

processing (including storage and memory) is instantiated by physical systems.[4]

When characterised in this way records are easily comparable to a measurement. Measurements are designed to gain information about a certain system (external to the measuring device) and present it in readable way. This connection has been picked up repeatedly in the literature (see Albert 2001, 2015; Wolpert 1992; Hemmo & Shenker 2012) and records are generally seen to be the outcome of a measurement (although not all measurements produce a lasting record).

Measuring devices are modelled in a simple way. The device starts in a known *ready* state. It then interacts with the system it is measuring and the state of the device is altered in a way that tells us something about the interaction. Similarly, Albert says that information is gained from records by a comparison between the current state of the system (the record) and a state of the system at another time (the ready state). This contrasts with retrodiction that uses only the current state and the dynamics. Using the ready state to calculate the evolution of the system, we can work out what the expected current state should be. Any deviations between this expected current state and the actual one indicates that an interaction with another system occurred in the interval between the time of the ready state and the current time, which altered the state of the record. The deviations in the actual state give information about this interaction. Albert uses a billiard ball scenario to exemplify the use of ready states: a ready state shows that the billiard ball was moving 10s ago and we should expect it to continue this motion (absent friction etc). But the present record is that the billiard ball is stationary. This allows us to infer that a collision between the billiard ball and another ball must have taken place in the last 10s to stop the first ball.

The major drawback of modelling records in this way is that it is overly idealised. The ready state in the past facilitates the inference made about the interaction occurring in the interval before the present state. But Albert's account only briefly considers the inference itself and what goes into making it. It is here that a high level of idealisation creeps in. In the context of Albert's work this is understandable. His focus is on how records can be proved reliable in the face of the *reversibility*

---

[4] See Maroney and Timpson (2018), Timpson (2013, chapter 2) for a full review.

problem, which argues that retrodiction using Boltzmann's arguments for why entropy increases should lead us to believe that entropy also increases towards the past. This contradicts the information of a low entropy past given to us by the multitude of records all around us (both natural and created by humans). As a result he focuses on the ready states and how they can be compared to the past hypothesis - his proposed solution to the reversibility problem. He calls this the "mother (as it were) of all ready conditions" (Albert, 2015) that sets up the conditions for ready states at the local scale. But he does not go into much more detail of how this global asymmetry leads to local asymmetry. To connect the two he argues that to use a record we require a second record of its ready state. However, this separate record will have its own ready state, and hence a third record of this ready state is required. This leads to an infinite regress that Albert argues can only be stopped by the past hypothesis. But it is unrealistic to assume that when we read records on a day to day basis we have this infinite regress in mind. Albert acknowledges this disconnect by saying that a "foggy" sense of this process is "hard-wired into the cognitive apparatus of any well-adapted biological species by means of a combination natural selection and every experience and explicit study and God knows what else" (2015, p. 39). The details of how we practically make inferences is not explored as it goes outside of the scope of Albert's project. Huggett (2023) raises similar concerns and goes into more detail about how Albert's account relates to local inferences. He emphasises that the past hypothesis is a precondition to records but not a direct part of our reasoning. Understanding records at a local level requires paying attention to the inferences made and not just the ready state.[5] Once we move out of the context of understanding how the record asymmetry fits into the wider thermodynamic asymmetry it is important to go back and work out the function of records more locally.

Albert's account assumes a simple isolated interaction between just two systems. In reality there will be many systems interacting, each changing the state in a different way. This does not rule out the method of using ready states but it makes it significantly more complicated. The different interactions may obscure

---

[5] Huggett also uses ideas about noise that are similar to what I will present in the next section, but they remain an operational concern and not as an essential part of a record as this chapter will argue.

and overwrite each other, making it difficult to separate out what changes to the record were the result of which interaction. Although some systems are more isolated than others – and many measurement devices designed for the purpose try to achieve as high a level of isolation as possible, see for example the measurements made at CERN – there is no way to completely prevent interactions with the many, overlapping systems that form the environment. In addition to these external interactions, Albert's account assumes that the system's evolution is perfectly stable and predictable. For most cases this is an excusable assumption but it is not one that can be made universally. Systems can be prone to random fluctuations or other instability that could make it hard to give any precise prediction of what the current state *should* be.

While we may still be able to identify that some sort of interaction has occurred we may not be able to say much more about it. This severely undermines the informativeness of records and indicates that something in addition to ready states is needed to describe why we have such accurate and easily readable records of the past. This is particularly relevant to naturally formed records that extend beyond what we might typically call a measurement. These sorts of records may follow the overall process of a measurement but they are not systems designed to be so and hence lack the deliberate isolation that most measuring devices achieve. Yet we have many examples of naturally formed records in the world that are clearly informative. A broken branch records an animal's passing and is identifiable (at least to an expert tracker) despite there being no mechanisms shielding it. More than just a ready state is needed to sort through interactions between complex and interconnected systems and allow for any inferences to be made.

Another worry with Albert's account is that it makes records highly anthropocentric (even though he clearly recognises that records are naturally occurring). What makes a record a record is an observer comparing the record to the expected current state by means of another record. A comparison can be made using any system and as such any system can be a record. Records always have a human component when they are read, but it is extremely counterintuitive to suggest that the asymmetry in information between the past and the future is entirely of human making. We seem rather to be identifying asymmetric structures in the world and interpreting information that is encoded in them already. The example of billiard balls makes this clear. While in the set up described it is possible to get information from the billiard ball at rest by making this comparison, this scenario seems contrived. In most cases of records, even if the ready state account applies when looked at in detail, the information seems to be there waiting for us and no active comparison is needed. A photograph is an example of this. We can think of the ready state as the undeveloped film of the camera. But when getting information from a photograph we never seem to make an active comparison with this ready state. Once again this account struggles to accurately model how we actually get information from records on a day to day basis and how records are formed spontaneously in the natural world independently of people actively making measurements.

4.3 Non-Idealised Records

I will now present a new account of records as correlations that are robust against noise. It may seem an obvious conclusion that records will be affected by noise, this is not revolutionary. What I want to argue is that noise is not an afterthought but an essential part of the definition. This will be made clearer in section 4.4, but first this section will explore what noise is and how it leads to the introduction of redundancy as an important feature of records. Section 4.3.1 gives a more precise definition of noise and looks at how robustness makes a correlation identifiable. Section 4.3.2 and 4.3.3. look at case studies to show how redundancy is used to protect against it and how this has theoretical, not just practical, implications.

4.3.1 Robustness Against Noise

In addition to laying out how outputs are correlated to inputs, measuring devices also do a large amount of extra work to ensure that these correlations are stable and reliable. Similarly, what makes a record informative is when there is one correlation that is clearly identifiable over and above other, spurious correlations. This ensures that there is a one to one correlation between the record and the interacting system such that changes in the record correspond to changes of that one system and no others. Without this condition the record would not tell us anything, as the present state of the record would be dependent on any number of different systems. To achieve this the correlation must be retained through the effects of many different interactions and one single interaction must be the deciding factor of what state the record is in, with other interactions producing only negligible changes. We can call such a correlation *robust*.

To characterise robustness more formally we can consider the concept of *noise* that is often used in communication and signalling. Noise is commonly defined as unwanted modifications to a useful signal; it obscures the information in the signal and makes it harder to read. If we think of records as equivalent to an informative signal from the past then noise can be any modifications to the record that are not the result of the interaction between the record and the system being recorded. Such modifications act to obscure the information the record contains.

Unwanted modifications to a *desired* signal makes noise a very anthropocentric concept (the same problem that Albert's account has). We do not want to define records with the idea that they are deliberately encoded, this would neglect all of the naturally occurring records that exist. Instead, we can replace *unwanted* with *irrelevant*. Irrelevant interactions and modifications are ones that do not significantly contribute to the dynamic evolution of the system or the final form of the record. This mirrors the use of noise in Dennett (1991) where he considers the existence of real and objective patterns arising in nature that we use to make everyday predictions..

> "Where utter patternlessness or randomness prevails, nothing is predictable. The success of folk-psychological prediction, like the success of any prediction, depends on there being some order or pattern in the world to exploit." (Dennett, 1991, p. 29)

Folk prediction, as opposed to rigorous mathematical prediction, relies on stable patterns that appear in nature, and it stands to reason to assume these patterns help us understand retrodiction as well. According to Dennett a real pattern is identified when it is simpler to describe a state in terms of a pattern with a certain amount of noise than it is to give information about every detail of the state individually.[6] The patterns are identifiable despite noise partially obscuring them

---

[6] *Simpler* is defined in terms of data compression when transmitting information about the

and their existence is independent of whether or not we are able to detect them with our current best apparatus. Beings with different sense organs than us may be able to detect different patterns. What matters is that the pattern provides a description that is simpler and more efficient. This is often formalised further by showing that the evolution of the pattern can be effectively modelled by simplified dynamics.

This use of noise actually goes beyond the idea of noise described above as simply interference in signalling. This extended use will be returned to in section 4.4. But for now it is enough to define noise by recognising an objective distinction between a pattern in the world (the record) and data that is extraneous and acts to obscure it (irrelevant modifications to the record). Characterising records in terms of noise picks out the robust correlations that exist objectively, regardless of our ability as people to read and interpret them.

Large interactions that significantly alter the record do not count as noise and if these occur then the original correlation can be disrupted, replaced, or destroyed. However, if the interactions are small (or are such that they do not significantly interfere with the record) then they can be counted as noise. What might be counted as noise with regards to one pattern would in other circumstances be the record itself. For example, cosmic microwave background radiation is a record of the big bang but also acts as noise in telescopic imaging. If a correlation is robust against noise then it is preserved over time and we are able to identify the relevant interaction from the background interactions that the system takes part in.[7] In these cases we are able to use the correlation to gain information.

Putting this together a record must be robustly correlated to a past event. Robustness makes the correlation stand out against the background noise and transforms it into something that is identifiable and informative. A non-robust correlation, even if it is correlated in some way to a past state, cannot be informative in this way because the correct correlation is indistinguishable from spurious correlations. Before considering this in more depth, it is worth looking at a particular feature of records that appears when accounting for noise: redundancy. I will consider

---

state.

[7] Another factor is that noise is cumulative and will build up to eventually erase records.

two case studies that show first, how redundancy is used, and second, how this is important for our theoretical use of records.

4.3.2 Using Redundancy: Error Correction Coding

The most common method for protecting against noise is *redundancy.* One need only look at modern computers to see that some form of reliable memory systems are possible on very small scales where noise has a significant impact. On top of this we have the development of quantum computing, including quantum memory. So how is this achieved? In part the systems are specifically designed to be as insensitive to noise as possible; however, it can never be completely eliminated. This is especially true in quantum systems with the addition of quantum noise (making noise one of the biggest challenges in achieving quantum computing). Instead, error correcting codes are used. The most common form of this is adding a redundancy of information where the information stored in the memory is duplicated; this can either be a directly replicated bit or be a function of many bits. The likelihood of both the original information and the redundant copies being affected by errors is low, so mismatches between them indicate that an error has occurred.

This can be done straightforwardly in classical computing; the quantum case is more challenging due to the no-cloning theory that prohibits duplication. But the idea of using multiple systems to encode the same information still holds. Shor (1995) showed that error correction can be achieved by spreading the information stored in one qubit over a number of entangled qubits. In essence, the states $|0\rangle$ and $|1\rangle$ are replaced by the entangled states like $|000\rangle$ and $|111\rangle$. We assume a bit flip error on just one qubit, which creates states like $|001\rangle$ or $|101\rangle$. An operation can distinguish between these different states and a correcting procedure applied (for more details see Shor 1995; Nielson & Chuang 2000). This operation transforms the state back to the original without disturbing it via measurement (so that the qubit can be used in further operations). This use of redundancy is slightly more indirect than the classical case, with the information preserved in entanglement relations rather than directly comparing redundant copies. However, the redundancy of qubits is still essential. Only with this can a robust correlation be achieved.

A single particle memory can never be a reliable source of information as its state can be changed by errors. We cannot distinguish between correct and incorrect cases. As a result, even in the instances where the system is correctly correlated, we cannot use it as a record because we cannot identify a correlation. There is no way to determine whether the state of the system is a result of the past state being accurately recorded or a result of noise. It is only when we can identify the correlation (i.e. the correlation becomes robust enough to have at least a strong likelihood of being correct) that it becomes informative. When redundancy is introduced, the accuracy and robustness of these memory systems is dependent on the system as a whole and the information is spread across multiple redundant parts of it. Only taken as part of the redundant set can it constitute a record with a robust and identifiable correlation to the past variable of the system it is meant to be a record of.[8] How many identical systems are needed to produce a reliable record will depend on the type of physical system being used and what other noise damping methods are in place. As the number of systems goes up the reliability of the record increases. The method is described above with 3 qubits for simplicity but in practice more will be needed for the necessary operations to work sufficiently – Shor's (1995) original quantum error correction code used 9 qubits and was capable of detecting an error in just one of the qubits. Building this into quantum algorithms is a significant technical challenge.

Redundancy is a good method to understand how to deliberately design microscopic record systems. But it can also be applied more generally to understand records at all scales. Take a (digital) photograph. A single blue pixel cannot tell us anything; it could be correctly correlated to the photo's subject or it could be affected by errors. But when the blue pixel is surrounded by many other blue pixels as part of a picture of the sky the other pixels act to verify the information in the single pixel. A single black pixel amongst this blue set could be identified as a speck of dust on the camera or some other noise. Another case of redundancy is in footprints in mud. There are multiple layers of redundancy here. First you might have a line of multiple footprints. Then within a single footprint you have redundant parts; only a fragment such as the shape of a heel or the outline of the toes is needed to identify

---

[8] We can label the entire set as a single record or say that each member is a record only when taken as part of the set. This difference has no particular significance.

it. Finally the outline of the footprint is made up of many minuscule marks, each on their own meaningless and indistinguishable from the uneven surface of the mud. These are not records. But once a redundancy of them has been collected, forming an identifiable fragment, it transforms into a record.

4.3.3 Redundancy in Theory: The One Molecule Gas Memory Device

Szilard's (1929) one molecule gas memory device and the surrounding literature makes clear how the use of redundancy can have theoretical implications and is not simply an operational concern. The dismissal of the effects of noise in memory devices has caused confusion and slowed progress in understanding thermodynamics. The analysis of records presented above can bring out two points about this literature: first that incorporating redundancy is necessary when theoretically modelling memory devices and second, a stronger claim that has still not been fully appreciated in this literature, that molecular scale records cannot exist in reversible thermodynamics. This fact should be taken into account when discussing the reversibility and irreversibility of thermodynamics in the context of intelligent agents.

This thought experiment (where a molecule is trapped in one side of a box by a partition and its position on either side encodes a 0 or 1 value) was introduced as part of the debate around Maxwell's demon and the possibility of an intelligent agent using a memory device to capitalise on fluctuations to violate the second law of thermodynamics. The idea being that a fluctuation to a higher energy state (e.g. a particle moving upwards due to Brownian motion) could be trapped and the increase in energy made permanent. The argument against this is that the act of measuring the particle to determine if it has fluctuated to higher energies - which is necessary to allow for the fluctuations to be exploited - would have a thermodynamic cost that cancels out the decrease in entropy. This process essentially involves operating a memory system that gives information about the particle. This has led to a long back and forth debate about the minimal thermodynamic cost of computation, with both Szilard's principle and Landauer's principle attempting to quantify this and work out where exactly in the process of recording and erasing information the cost is paid. There have been extensive attempts (e.g. Ladyman et.al. 2007; Ladyman et.al. 2008; Ladyman & Robertson 2014) to prove Landauer's principle in particular, which have been objected to by Norton (2011; 2013; 2017a) as failing to account for fluctuations (noise) that make reversible thermodynamic processes at this scale impossible. The fluctuations that are being capitalised on also prevent the memory device from working reliably (e.g. the partition that attempts to trap the molecule in one half of the box will itself fluctuate). Norton (2017b) lays out the history of Szilard's thought experiment and the confusion that its inappropriate use has caused in responding to Maxwell's demon. I will not repeat the full analysis here and refer readers to the original literature for details - it is enough to say that a proper account of records would have assisted this literature and clarified the debate from the start.

This is not to say that the literature has not been productive and more recent literature on Landauer's principle has begun to address the fluctuation problem. It is easy to see that the methods used to do so are essentially introducing redundancy just as is used in error correction coding. Myrvold (2021b) builds on the earlier proofs mentioned above to prove Landauer's principle accounting for fluctuations; he also shows how these can be adapted for a quantum mechanical treatment. He uses expectation values, calculated from the probability distribution of the system due to thermal fluctuations, rather than specific values (more on this in section 4.5). As a result of using expectation values and probability distributions Landauer's principle holds as a statistical limit rather than an absolute one.

Redundancy can be identified in the use of expectation values and probability distributions. These represent the statistical result of a long sequence of trials or a set of identical memory devices all undergoing the same procedures. Statistically the majority will be unaffected by noise and will be correctly correlated. None of the individual systems are informative but the ensemble as a whole is robust against noise and therefore can act as a reliable record. Incorporating this into the theoretical treatment has allowed the literature to move forwards.

However, this analysis has limitations and we can question what exactly has been proved. Although it has established an undeniable link between logical and thermodynamic irreversibility, it is not clear that is has established anything about memory devices in particular. Myrvold and Norton (2023) have recently clarified that Myrvold's analysis only accounts for fluctuations in the system itself and not the apparatus used to control it. (In his account the particle is trapped by an external potential, fluctuations of the molecule within the potential are accounted for but the external potential is assumed not to fluctuate and manipulating it is not assumed to take place in a reversible way). This isolates the logical operation taking place in the system that takes two (approximately) distinguishable states to a single state (a logically irreversible many-to-one operation). But, although having two distinguishable states is a basic requirement for records, this is not a sufficient condition. It must also be reliably correlated to something; this is necessary so that the information encoded is not random data but a meaningful record that allows inferences about the past. The procedures that create such a robust correlation are exactly those that Myrvold's analysis still neglects. The apparatus must also be fluctuation free for it to have the necessary control to carry out the procedures in a slow and stable enough manner to satisfy thermodynamic reversibility (in other words we have not ensured that there is no dissipation of energy from the system). The fluctuations inherent to any actual processes means a record cannot actually be created in a thermodynamically reversible way. They write that:

> "Myrvold...seeks to place the thermodynamic cost computed in Landauer's principle on a firmer theoretical foundation by deriving it within statistical physics, without assuming as a primitive the precarious notion of a thermodynamically reversible process. The analysis deals with systems subjected to external manipulations, represented as changes to the system's Hamiltonian, and focuses on dissipation within those systems, setting aside entropy generated by the external systems driving the changes. It is argued that, in addition to any entropy generated by the external drivers, there is entropy generated within the manipulated system that is subject to the Landauer bound." (Norton & Myrvold 2023)

Hence for proving Landauer's principle as a minimal cost to logical operations Myrvold's analysis is sufficient. This literature does not depend on its being a reliable record, only on the distinguishable states discussed above. However, in the context of Maxwell's demon and the actual abilities of an intelligent agent to physically violate the second law, this is significant. The idea that this is a reliable memory device is central, as is the requirement that this is all done in a reversible manner with no dissipation both from the system and from the apparatus used. The concept of records has been incorrectly assumed to hold unproblematically in this domain and the account of records presented here shows why this is impossible; the necessary robustness cannot be achieved. Introducing redundancy solves some of the issues but a full treatment in these terms will take us outside of the realm of reversibility. Norton's analysis indicates to this conclusion but does not emphasise it, nor does it make clear that the one molecule gas memory device is *not* a record at all rather than an improperly functioning one. So an intelligent agent cannot manipulate a system in a reversible way to take advantage of fluctuations to violate the second law because the tools needed to do so (records) are not definable under these conditions.

4.4 Robustness and Stability

Having established a definition of noise and the requirement for records to be robust against it, I will now develop this idea further to show how noise is definitional of records. So far I have mainly considered the simple sense of noise similar to how it is thought of in signalling. However, the way noise is defined is more general than this and can also help us understand how many systems in the world naturally provide robustness. The world is abundant in stable and easily predictable (and retrodictable) systems; this sort of general stability, which can be characterised in terms of noise, plays an important role in records alongside redundancy and explains how robust correlations can be created in non-isolated systems. Stability also defines the difference between records and retrodiction, a distinction which is made use of in the literature on time asymmetry.

4.4.1 Stability and Inferences

So how does stability figure into an account of records? It is a frequent precondition of robustness that acts to single out a single correlation from background noise. Effectively it defines semi-isolated systems in the world that are only sensitive to certain types of effects. To get to this I will first consider how stability plays into our inferences before looking explicitly at how it singles out a single robust correlation through a comparison with measuring devices.

Stability can be used to expand on the method of using ready states to make inferences about the past. In fact, in many cases we can do away with the use of ready states entirely and instead rely totally on stability.[9] As previously mentioned, Albert uses the billiard ball scenario to exemplify the use of ready states: the ready state shows that the billiard ball was moving 10s ago (which we know through a separate record and use to form the expectation that it is still moving now) and the present record is that the billiard ball is stationary. In this case a separately recorded ready state is necessary to be able to make any inferences, but in many cases we can make exactly this sort of inference without knowledge of any specific ready state. When systems are reasonably stable and easily predictable we already have extensive knowledge of different types of systems and the typical evolution they display. Using this we are already able to make fairly accurate informal inferences about what the past was like. And from this infer the expected present state (absent interactions), hence doing away with the need for a ready state. Most of the time it is not even any specific state that we make guesses about but rather just what the typical evolution of a system would be and seeing a deviation from this indicates an interaction in the past. Consider a broken branch. Our knowledge of the laws of both biology and physics governing branches indicate that this is not the typical evolution of an undisturbed branch (a branch will not break itself) and so we can immediately see that there must have been some interaction that broke the branch. This inference did not need any knowledge of a specific time in the past where the branch was unbroken but relied on the general stability of the system and our expectations of it.

The role that stability plays in inferences feeds into the distinction between records and retrodiction. The main difference between the two is that records have a direct correlation to a specific past state and we can largely ignore the evolution of the system between the time of interaction (that created the correlation) and

---

[9] This is not to say that Albert's account is not helpful in any circumstances, merely that ready states are not necessary and often not used at all.

the present, and we can ignore completely the evolution of all systems except the record. Retrodiction, meanwhile, evolves the present state backwards so this evolution is essential and must include all relevant systems. It can also take us to any arbitrary past time and not a specific event. Albert characterises this as a sharp distinction between records and retrodiction; describing them as fundamentally different methods for getting information about the past.

However, looking at retrodiction in a more practical light reveals that this sharp distinction is misleading. Retrodiction in its most basic form is the process of taking the present state of the world and evolving it backwards using dynamical equations to find what the state was in the past. Full retrodiction would require taking the exact microstate (if such a thing can be said to exist) of the whole universe – or at least an area large enough to cover all possible influences on the system we are trying to predict. But in practice there are computational limits that make this highly impractical. Instead we tend to use an idealized, isolated subsystem that can still produce accurate results. This involves many techniques such as making assumptions or approximations about microscopic details and neglecting dynamical contributions judged to be irrelevant. On top of this, we can often use our knowledge of a system, the laws governing it, and causal relations involved to make inferences about a likely past state with no (or almost no) calculation at all. Fallen leaves under a tree allow us to infer leaves on the tree above as a past state. A ball in mid-air allows us to infer the trajectory it took to get there. But this is still essentially no different from the retrodiction method. All these inferences can be made because we have knowledge of the world and its typical dynamics. We can make quick and simple inferences rather than doing precise mathematical calculations because there are many stable systems in the universe that are not dependent on the exact microscopic state of the world, especially not the state of the world outside of the system we are interested in. We can take advantage of these stable patterns of evolution to ignore extraneous details and skip doing the calculation each time.

A more rigorous version of this inference technique is to make a guess at a plausible past state using the shortcuts described and evolve it forwards using dynamical laws to verify it against the actual present state. Wallace (2017) describes such a "guess and check method” that he calls historical inference.[10]

Initially, records are clearly distinguished from the process of retrodiction. For example consider what would happen if we wanted to find out what the night sky looked like two weeks ago. Retrodiction would mean taking the present state of the night sky and evolving it backwards using the laws governing the movement of stars across the sky, considering the rotation of the earth etc, to find the position that each and every star and planet would have been in (frankly this already requires a large number of idealisations and shortcuts based on astronomical models). But if we have a photo of the sky two weeks ago we can get this information instantly just by looking at it without any sort of calculation. This direct informativeness is a defining feature of records. But when we consider more common and realistic examples of retrodiction using shortcuts based on idealisations and stable systems this difference becomes far less pronounced. We can jump immediately from fallen leaves on the ground to previously green leaves on the tree. The method by which records give information is not dramatically different from retrodiction. Starting from a complete retrodiction we introduce more and more shortcuts until we are able to make inferences towards the past in the direct and instantaneous way that records allow for.[11] Techniques such as historical inference and methods using causal dependencies sit somewhere along this path.[12] And all of these methods

---

[10] Wallace also claims a sharp distinction between retrodiction and historical inference that is no more realistic than the distinction between records and retrodiction.

[11] Although records and retrodiction are closely connected it is the multitude of memory systems that leads to the asymmetry between the past and the future. Retrodiction is matched by prediction but records disproportionately give information towards the past. The shortcuts needed to get from a full mathematical prediction/retrodiction to directly informative records are biased towards the past. On top of this a single record (which gives direct information without any regard to the intervening evolution of a system) can be used as a basis for many more inferences about the past beyond what is directly recorded. An example of this would be animal tracks on snow. The physical print in the snow acts as a memory of a paw shaped thing impacting on the snow and causal inference allows us to attribute this paw shaped thing to an animal walking. We can also make generally accurate guesses as to what causes are. For example we generally rule out the possibility that paw prints on snow were caused by a hovering aircraft with paw shaped sticks being used to create the tracks because knowledge of the world makes the possibility of an animal walking far more likely an explanation than this outlandish possibility. Hence records lead to a dramatic asymmetry in what we can know about the past versus what we can know about the future. The existence of paw prints in the first place is what allows us to make these inferences about what created them, but a blank field of snow will never allow us to make guesses about future animals walking across.

[12] The various accounts of records do not seem to agree on what exactly counts as a record

are commonly used to furnish the background information used to replace Albert's notion of a ready state.

One might be concerned here about how eroding the distinction between records and retrodiction is at odds with accounts of time asymmetry such as Albert's that seem to rely on the distinction. The reason for the distinction was to allow Albert to use records to justify introducing a past hypothesis to counteract the reversibility problem. The reversibility problem occurs from time reversing the methods used to derive entropy increase in the future. There is no inherent asymmetry in the methods so applying them backwards retrodicts entropy increase towards the past as well, leading to the conclusion that we are Boltzmann brains that have fluctuated into existence at this very moment. The distinction of records from retrodiction, and the direct connection between records and the past hypothesis that Albert argues for, exonerates records from this concern. If records rely heavily on retrodiction then we cannot use them in this way.

This worry can be mitigated to some extent. Albert draws his distinction specifically using retrodiction within statistical mechanics. However, the way that I have considered retrodiction is much more general and uses all possible scientific theories (many of which do not have even a pretence of symmetry). So the context is different and problems with retrodiction are not found in all domains.

Additionally there is a general concern about potential explanatory circularity in using empirical evidence to justify the choice of laws and initial conditions and then using these laws to explain the evidence. This problem is already a topic of discussion in the literature on laws and the past hypothesis. Eroding the distinction for records may add to that particular problem but does not create or redefine it.[13]

---

or not. There are systems that are clearly records – i.e. photos, computer memory, etc – and systems that clearly are not and can only give information through retrodiction. Albert's account is best suited to clear cut cases of records and as a result makes an erroneously sharp distinction between records and retrodiction. Reichenbach's account or the fork asymmetry as considered by Stradis (2021) would likely count historical inferences and causal connections as types of records. Reichenbach defines records just as subsystems with low entropy compared to their surroundings. Reichenbach, then, is much more permissive on what counts as a record. Where we start applying the label records has no real significance and is reflective of the close connection between records and retrodiction.

[13] Uffink (2002) criticises how Albert confusingly conflates and uses the past hypothesis, records, and empirical evidence.

4.4.2 Stability and Noise

Looking again at how records are related to measuring devices finishes off this analysis and shows how stability singles out robust correlations from noise. This takes us the final step from retrodiction to records. Increasing stability makes retrodiction easier and easier and eventually allows for the necessary robustness that can single out a direct correlation (robustness of the system and of the correlation are interdependent). A measuring device requires three conditions be met:

1. Specific types of inputs interact with the system in a way that produces an identifiable change.
2. Different inputs result in different outputs (different identifiable changes to the measuring system).
3. The system is not (significantly) further altered by anything after the desired interaction.

The first condition is what creates the correlation between the record and the system. The second ensures that the two are genuinely connected and that the correlation is not due to some fluke, conspiracy or fine-tuning. The final condition is necessary so that the correlation is preserved over time and the record does not become distorted or unreadable.

The direct correlations between inputs and outputs used in measurements allow us to ignore the evolution of the system; this is what makes them informative. However, this is slightly misleading and is a result of an idealised understanding of a measuring device. Although when using a measuring device we ignore the evolution of the system being measured, the evolution of the measuring device system is of critical importance. This is particularly clear when we look at designing a measuring device. To set up a device that satisfies the above criteria we examine a system in great detail until we can predict exactly what will happen to it both when left on its own and when it is allowed to interact with a variety of different external systems. Once we have this information we can fine tune the system to only react to the type of inputs we want to measure and can generally ensure that the system works in the way we want it to. This process is effectively making the device robust in all the ways that have been described so far. Only once this is done can we forget about the evolution of the system and simply use the established correlations between input and output as informative correlations. The measuring devices designed and used by humans tend to be very specific and complicated systems, and hence ready states are often necessary at first because without them the system is not understood in enough detail to easily spot deviations from the expected evolution. Knowledge of a specific ready state, as in the billiard ball case, allows to us to create an ad hoc measurement device and turns something that would not ordinarily be a record (such as a ball stationary on a table) into something informative. But many records, and particularly naturally formed records, rely on the sort of stable and easily predictable systems for which a ready state is not needed to work out the necessary information.

Stable systems are objective features of the world. However, our ability to identify and use stability may vary. It is dependent on our current best technology and understanding of science as well as the knowledge of each individual. Going back to the distinction between records and retrodiction we can now see that how directly we get information from correlations depends on our understanding of the system. For example a scuff mark or a broken branch are a record for an expert tracker because they understand the systems involved. But for the average tourist these would be meaningless and any information would require doing some sort of complex retrodiction.

Finally, this aspect of records can also be understood in terms of robustness against noise: stability implies that a system is not susceptible to noise. So far we have mainly considered noise to be the result of external influences such as surrounding electromagnetic fields, light pollution etc. But it can also come from internal factors in the system such as thermal movement and electrical fluctuations. On top of this there are many small details of the evolution of a system stemming from small differences in initial conditions that are generally irrelevant to the overall evolution of the system. All these factors together influence the exact evolution of a system; even in a perfectly isolated system the internal sources of noise must be considered. A system is stable and easily predictable when the overall evolution is independent of these sorts of details. The same system can be seen many times in different environments with slightly different initial conditions but will still act mostly the same because its evolution is robust against these sources of noise. Due to its robustness the system will display the same behaviour on different occasions and so provide a basis for making easy inferences about the past. It is not a coincidence that these are the sorts of stable systems that are likely to be involved in the robust correlations that are required for records. This use of noise relates strongly back to Dennett's original use of the concept as discussed in section 4.3.1, which goes beyond noise as it is regularly thought of in signalling theory. These sorts of stable patterns that are describable under their own simplified dynamics are exactly the sorts of patterns he is looking at.

4.5 Robust Records must be treated Macroscopically

What use is this account of records? Section 4.3.3 already showed that it denies the possibility of microscopic records in reversible thermodynamics. This generalises into a wider implication that is worth exploring: records should be treated using the sort of methods we use to handle macroscopic physics. I will not go so far as to claim that records must all *be* macroscopic. What counts as macroscopic is vaguely defined and we can certainly make our records very small (the case studies in section 4.3 are examples of this; even when redundancy is introduced the resulting systems are very small). Instead, the scale of records is linked to their reliability and how informative they are. Microscopic records are fleeting, hard to identify, and normally carry little information. When a human designer is introduced we are able to make records very small by shielding the systems involved as far as possible, but there are limits to the extent this can be achieved.

My account makes it clear why records, especially naturally occurring ones, are *likely* to be macroscopic. The main criterion - of being robust against noise - is most easily achieved by stable, macroscopic systems. It is well known that we can have macroscopic dynamical laws that accurately describe the evolution of systems while ignoring the microscopic details. We are also able to treat such systems as effectively isolated from distant effects and small interactions that might occur.[14] The need for redundancy in records also increases the size of the system needed

---

[14] Elga (2005) lays out some reasons for why the macroscopic world has these features

and pushes us towards the macroscopic domain. The likelihood for records to be macroscopic fits into the project of explaining record asymmetry. It is commonly found when looking for asymmetry in physics that it is missing from the fundamental level but emerges at a higher level. Records are no exception to this. They are frequently referenced as being observable (linked to their epistemic availability) or otherwise as macroscopic. Yet no justification of why records specifically should be macroscopic has been given so far. This account answers that question.

But beyond a likelihood of being macroscopic, I will also make the stronger claim that when treating records theoretically we must adopt the sort of techniques - for example defining macrostates, using statistical sets, or coarse-graining - that are abundantly used when developing macroscopic theories and discussing their emergence from lower level theories.[15] This is exactly what was done for the one molecule gas memory device to accommodate the need for redundancy. Explicitly incorporating stability will require similar methods. Recognising the role these techniques play in records has the potential for numerous explanatory benefits both for the asymmetry project and more generally. Accounts such as Albert's take records to be connected to but separate from other asymmetries such as the thermodynamic asymmetry. But properly modelling records using these techniques reveals how intertwined records are with many other theories and with macro-micro relations more generally. This should change the way that we think about records and their connections with other questions.

I will sketch one such implication here: what thinking about coarse-graining and records could mean for explanations of time asymmetry. Understanding records in terms of coarse-graining could provide a stronger argument for their asymmetry than Albert's account can. Albert claims that the past hypothesis is a type of ready state and this is how records can be directly linked to it. However, his past hypothesis, which places only a general constraint on the initial state (that it is low entropy), looks nothing like the sort of ready states he uses that specify an

---

[15] For example, one place where thinking about macroscopic techniques while modelling records could have explanatory value is quantum decoherence, in the accounts of quantum darwinism and decoherent histories. The next chapter will look into this. In the decoherent histories account records mark when sets of histories are consistent and are often tied to the coarse-graining used to derive classical dynamics (see Gell-Mann & Hartle 1993).

exact state and are used to predict the exact expected present state. The general similarity in techniques is only that both ready states and the past hypothesis create a sandwich effect where knowledge of the present and a point in the past restrict what could have happened in between. Beyond this the connection between ready states and the past hypothesis is unclear and the explanation of asymmetry seems to falter.[16] Meanwhile, Wallace's (2012) version of the past hypothesis aims to explain the general success of asymmetric coarse-grained macro-dynamics. If records can be connected to coarse-graining techniques an explanation of their asymmetry and its connection to the past hypothesis would fall out naturally as part of this broader picture.

A more formal exploration of how records work with coarse-graining is difficult for a number of reasons. Most work on coarse-graining (including Wallace's) focuses on clearly defined dynamics that records do not obviously have. Case studies such as records in decoherence, which will be the subject of the next chapter, and the case studies considered earlier may help provide a firm structure. Going back to Myrvold's formulation of the one molecule gas memory device, the methods he uses of probability distributions and expectation values are potential examples of coarse-graining that are very similar to the techniques used in defining entropy and deriving thermodynamics from statistical mechanics. However, whether or not they fit the definition of coarse-graining is a complicated issue and what counts as coarse graining varies across the literature. I will not try to answer this question here. Additionally, even if this can be worked out in particular cases it will not obviously generalise. Records systems can be formed out of potentially any type of system and a common structure cannot be assumed.

However, despite these uncertainties, considering these sorts of techniques provides a valuable route for exploration and gives a much stronger foundation for how to treat records when we encounter them in our physical theories. Understanding how redundancy and stability require the use of methods such as statistical sets, coarse-graining, or just the general use of macroscopic theories, serves as a starting point to understand why and how records are being used throughout physics.

---

[16] Frisch (2023) develops a similar objection to Albert in more detail.

4.6 Conclusion

In summary, records can be defined as systems that enter into robust correlations with the past state of some other system. The robustness criterion is the key feature that leads to the informativeness of the record and allows us to identify the information encoded in the system. The second conclusion of this chapter is that robust correlations require stable macroscopic systems and redundancy to provide background information and protect against disruption from noise. This description of records explains how they are used day to day and takes into account the practical limitations they face, as well as highlighting why considering these features is so important for the way we treat records theoretically. It establishes them as objective features of the world that we are able to access and interpret. Applying this account of records to instances where they are used in both physics and philosophy will expand our understanding of how information is being encoded in the relevant processes and what role records are playing.

By understanding records in this way it has also become clearer what is required to answer the question of why we have a record asymmetry and how this relates to time asymmetry. By making explicit the reliance of records on stable, macroscopic laws from physics, biology and other high level sciences, some parts of this can be answered already by looking at the work done on the emergence of an asymmetric macroscopic world from a symmetric microscopic one. The background asymmetry feeds into the asymmetry of records. But it remains to be examined exactly why robust correlations seem to form between the past and the present but not the future and the present. A further exploration into coarse-graining and other such techniques as suggested here could provide this answer.

But even pending a full explanation of this, it is undeniable that records *do* play an important role in the emergence of asymmetry in statistical mechanics. And with a clear sense of what records are and what features they have we can clearly incorporate them into our model of time. Robust records are distinctive features of reality; they are not arbitrarily picked out or reliant on human understanding. As such, records provide a firm structure which can be used to ground truth claims about the past, and do so in a way that naturally connects this metaphysical asymmetry to the physics of emergent time asymmetry.[17]

---

[17] A more detailed analysis of how records can act as truthmakers I take to be outside of the scope of this thesis. Roughly speaking we can take Wilson's (2013) account of indeterminateness as a guide where the state of affairs (i.e. a present state containing robust correlations) is the determinate for a determinable. But there may be complications relating to how certain the inferences made from records are - after all records allow us to infer with a certain degree of certainty but do not give absolute guarantees. There are also more general concerns about what it means to be a truthmaker, how to turn this into a semantic theory of truth, and how this changes when we have to think about emergence (see Barnes (2012) for some discussion of how truth relates to emergence). These questions remain unanswered; but the physical basis, and the justification for it, have been made clear.

5 The Role of Records in Quantum Decoherence

5.1 Introduction

Another area which has made extensive use of the idea of records is quantum decoherence. I will now turn my attention to how records are used there, in particular in the theory of quantum darwinism. There are three aims to this: 1) to further explore the important role that records play in physics and to showcase how the account I developed in the previous chapter can be used to strengthen our interpretations of theories, 2) to prove that information can be used to understand a formal theory of emergence (I have already shown that records are predominantly macroscopic but I have not looked at the idea of emergence more generally), and 3) to show how quantum mechanics fits into the model of time I have developed. The previous chapter, and the majority of chapter 3, focused on how probabilities and records in statistical mechanics are the basis for the open future and the fixed past (alongside other factors such as the lightcone structure of spacetime). I argued that a reduction to fundamental probabilities - such as the ones in quantum mechanics[1] - is not necessary and that we should accept an emergent account of time's flow. However, despite discrediting quantum mechanics as the sole option for finding the open future in physics, it is still a contributing factor; looking at decoherence specifically will highlight how quantum mechanics is connected to emergence and is continuous with higher level theories such as classical mechanics.

To achieve this broad set of goals I will focus here on the theory of quantum darwinism - an account of decoherence that aims to explain the emergence of classicality in information based terms by analysing the production of records during interaction with the environment. This will go together with the following chapter that will continue this project by looking at an alternative interpretation of quantum mechanics (relational quantum mechanics) that focuses more on the perspectival element of this project. Together these two chapters will demonstrate how an information based, emergent, and perspectival approach to quantum mechanics is a

---

[1] Caveated that quantum mechanics itself is not necessarily fundamental.

robust and widely applicable method of interpretation. I will not comprehensively defend either quantum darwinism or relational quantum mechanics but instead look at how they support the model of time I am developing and show in turn that thinking about time in this way can assist these theories.

Specifically in this chapter, I will show that the account of records developed in the previous chapter strengthens quantum darwinism by providing a clear standard for what physical features the world needs to have for observers to be able to operate in a way that is familiar to us. Additionally I will argue that this standard can be used to connect quantum darwinism to philosophical accounts of emergence which use the screening off of lower level details. The term emergence is used loosely in physics but philosophically far more critical attention has been given to what exactly constitutes emergence; precise conditions such as novelty and autonomy have been set. Quantum darwinism tells a conceptual story about how the ability for multiple observers to independently verify a quantum state is the "hallmark of classicality" (Zurek 2009, p.182); it then attempts to precisify this using limits in information theory. However, the conceptual aim for the project remains vague (and largely disconnected from the formal work) and little attention has been given to understanding it; the standard of classicality that quantum darwinism aims to derive is unclear. This is in contrast to other accounts of decoherence such as decoherent histories that definitively derive classical dynamics out of quantum mechanics. This standard - what it is and why it it distinctly classical - sets a clear goal and is used in the philosophical literature to justify emergence through the novelty and autonomy of these dynamics (e.g. Wallace 2012; Franklin 2022). The lack of a corresponding standard for quantum darwinism leaves it hard to understand what should be used to assess the conditions for emergence. It has also contributed to a number of criticisms of quantum darwinism regarding its use of observers.

I will first lay out quantum darwinism and some criticisms against it, as well as looking at how it compares to the decoherent histories program using a model of scattering photons off a central atom. Then I will explain how records - understood as correlations that are robust against noise - can provide a standard for classicality; the use of redundancy in particular will be essential to this analysis. Finally I will look at how my account of records can connect quantum darwinism to philosophical accounts of emergence, and in doing so show how information can be used to understand the relation between levels. This addresses the first two aims that I laid out above, I will return to the third goal (showing how this fits into the model of time I am developing) in the conclusion by looking at how quantum darwinism connects to and extends the already established literature on decoherent histories that argues that records are used to 'fix the past'.

5.2 Quantum Darwinism and the Emergence of Classicality

5.2.1 Quantum Darwinism: A Review

Environment-induced decoherence generally proceeds as follows: a system interacts with the wider environment around it and becomes entangled with it. This leads to a suppression of coherence in a particular basis of the original system. The original approach to modelling this was to model the environment in some way, allow it to interact, and then trace it out to leave the state of the system of interest. The reduced density operator representing the state is transformed into an effectively diagonalised form; this is interpreted as being a classical state as it approximately obeys classical probability calculus and its evolution follows classical dynamics (for details on how this works and why the resulting state can be interpreted as classical see Zurek 1982, Schlosshauer 2007, Wallace 2012 - note again here that this does not necessarily claim to solve the measurement problem in its entirety). The environment is a means to an end in this sense. Used and then discarded.

This program developed in connection with the parallel work on consistent histories by showing the equivalence of distinguishable histories and the suppression of interference (see Rocha, Rickles, & Boge 2021 for a historical picture). Most modern presentations of decoherence (such as those mentioned above) start with the diagonalization of the density matrix and then present the consistent/decoherent histories approach to model the dynamics over a sequence of times. The latter framework does not (supposedly) require an open system so has been better suited for work in, for example, cosmology (Gell-Mann & Hartle 1993), as well as helping with foundational issues such as the measurement problem (see Okon & Sudarsky

2016 for more on this as well as criticisms).

Gell-Mann and Hartle did significant work on decoherent histories (as consistent histories came to be called) and identified that the diagonalization of the density matrix was only a weak condition for decoherence. They focused on incorporating the additional effects of noise, fluctuations and dissipation to produce stable classical dynamics (Gell-Mann & Hartle 1993; 1995). One of the new criteria they introduce was for histories to be distinguishable by generalised records - a set of orthogonal projection operators on the set of histories. These ‘records’ distinguish between possible pasts, but are called generalised as Gell-Mann and Hartle note that they are global projectors and do not have the common properties of records, properties like being preserved over time (although their work on nested records starts to build on this) or being accessible to observers. Instead they are global markers for which history is which.

Now we get to quantum darwinism, which largely developed out of Zurek’s early work on decoherence. Quantum darwinism puts records at the centre and can be seen as a more practical account of the idea that records play a role in decoherence. An explicit connection to consistent histories (and the work of Gell-Mann and Hartle) is shown in Riedel, Zurek and Zwolak (2016); but largely quantum darwinism focuses on how records - with all the common features associated with them such as observability - are produced in a single case of decoherence (rather than at a sequence of times like in decoherent histories).

Instead of looking at the dynamics and behaviour of the decohered system, quantum darwinism instead focuses on how a defining feature of classicality is the possibility for objective agreement among observers. Something made impossible in quantum mechanics due to measurement affecting the quantum state. Either due to the collapse of a superposition; or more weakly, if we reject collapse, because once measured the systems are entangled and cannot be viewed separately. Hence a second observer cannot independently verify the state of the system with no reference to the first measurement. Quantum darwinism argues that classicality - in terms of intersubjectivity and the ability to verify measurement results - can be achieved because decoherence creates many redundant records of the state of system (in the pointer basis) in the environment. These records allow for indirect measurements that have minimal effect on the system itself, therefore allowing multiple observers to act independently. Many observers can measure different fragments of the environment and separately get information about the state. As such tracing out the environment - as is traditionally done in environment-induced decoherence - misses important details. The key to this is redundancy that ensures that each fragment will contain informative records. This explains how this element of classicality - the possibility of objective, shared information - is achieved.

It should be noted that decoherence on its own is not generally taken to completely solve the measurement problem; most frequently it is paired with an Everettian style interpretation of the wavefunction and decoherence is taken to explain why each branch *appears* classical. Within each branch there is an effective wave-packet collapse leaving the appearance of robust classical states and outcomes that can be predicted by the Born rule. Despite its focus on observers, quantum darwinism should be seen in the same vein; it is identifying a standard for what it means for the world to appear classical but not claiming to derive absolutely how an observer gets a single outcome out of unitary evolution.[2]

The mechanism that creates records is also supposed to explain how the pointer basis is selected. The no-cloning theorem prevents the copying of unknown quantum states so this cannot be what is happening. Instead what *can* be copied is observables (sets of orthogonal states). The structure of the environment, and how it interacts with the system, determines which observable will be most effectively copied and hence selects this as a pointer basis. The copying reinforces this observable and suppresses other possible bases. This is the motivation for the name *darwinism*: states that reproduce themselves most effectively are the fittest for survival.[3]

---

[2] I will largely avoid discussions of interpretation in this chapter but I will return to some discussion of this at the end. Zurek's writings also contain extensive discussion of his own interpretation - the existential interpretation - but this goes beyond just the formal theory of quantum darwinism and is not a necessary consequence.

[3] This mechanism of interaction with the environment both selects the pointer basis and creates records. However, it is possible to have decoherence (defined as basis selection - this corresponds to the effective diagonalization of the density matrix) without records being created. This is possible if the environment is not suitable to store information in the forms of records; for example the records are immediately disrupted due to noise. Hence quantum darwinism always implies that

Quantum darwinism can be demonstrated in a practical model: scattered photons. Photons scattered off a central atom in a position superposition each encode information about the decohered position. Roughly speaking the photons act as records. This is formalised using the concept of mutual information

$$I_{S:F} = H_S + H_F - H_{S,F} \tag{1}$$

where H is the von Neumann entropy and S and F denote the system and a fragment of the environment respectively. This measure (roughly) represents how much information can be extracted from one system about the other; in other words how much information can be extracted from the fragment of the environment - in this case from a certain number of photons - about the target system. Mutual information is especially useful in the context of decoherence as it can capture the difference between classical and quantum information. In the classical realm knowing the state of a composite system implies knowledge of its parts. This no longer holds true for quantum systems where the state of the whole cannot tell us about the parts; the parts do not have well defined states at all. Measurements on a joint quantum state reveal more information than any sequence of measurements on its component parts. The classical information contained in a fragment of the environment has a mutual information equal to the entropy of the system H_S. The complete quantum information about a system goes beyond this, however, as it contains more than just information about the composite parts. Hence total information about a quantum system goes to $2H_S$.

Riedel and Zurek (2010) model how mutual information depends on the size of the fragment measured and on the elapsed time (time during which the scattering process progresses and the number of scattered photons increases). They use photons that are not energetic enough to individually resolve the superposition. They find that during short time periods where few photons have been scattered mutual information i.e. the information we can get from a fragment of the environment about S, remains low. This indicates that decoherence effects are minimal and that

---

the system is decohered, but a decohered system with a diagonalized density matrix does not necessarily imply redundant records (section 5.4.1 explores this further).

the photons do not tell us about the position of the atom. However, for slightly longer time periods the information gained from a small fragment quickly reaches the classical value H_S. Once this is limit is reached, measuring larger environment fragments does not lead to information beyond H_S. The only way to retrieve the total information $2H_S$ about the system is to measure the entire combined system and environment.[4] This plateau around the classical information value corresponds to the atom being decohered and the information about its state being redundantly recorded in the environment.

The production of records therefore creates a particular kind of environment state that encodes the classical subset of the total information available. The plateau around classical information is not a general feature of an arbitrary quantum state. It is an identifying feature of decoherent states. This, then, defines the emergent classical state that we are looking for. It obeys the rules and limits of classical information with only negligible contributions from the quantum state that affect the state only when the entire system is measured or when the system has interacted for very short time periods.

5.2.2 Criticisms: Explaining Observers

This goes some way towards satisfying the question raised in the introduction: what standard of classicality is quantum darwinism aiming to derive? Classical information theory sets precise limits that can be recovered analogously to how classical dynamics can be derived. However, the formulation in terms of information theory is very different from the motivating statements about observers. There is a generic way in which information theory describes how information can be extracted and manipulated by agents; but when we consider observers more closely we do not tend to formulate it in these terms. Instead we think more practically in terms of measuring devices and inferential practices. It is not straightforward to use information theory to capture this much more general sense of information and observers; there is little analysis being done to pin down how the limits in information theory correspond to the everyday assumptions that observers make

---

[4] In a many worlds framework this would require the measurement to include the different branches. Hence the observer would need to be outside of the decohered branches themselves to make this measurement.

or the exact constraints they operate under. There is disagreement about the exact physical meaning of mutual information (see Korbicz 2021) and its definition remains technical rather than conceptual. Clarifying a standard for classicality that is more explicitly connected to the conceptual motivations is necessary if we want to understand why the resulting theory explains the classical world that we experience. Once this is in place can we start to look at how the limits in information theory represent this conceptual standard (although I will not attempt this here).

The confusion over how exactly quantum darwinism explains observers is revealed in the criticisms that have been levelled against it. Fields (2010; 2011) argues that there are several classical assumptions observers have to make before they can use the records in the environment to determine the state of the systems. These assumptions, he argues, show that quantum darwinism has not explained how agreement among observers arises *purely from quantum mechanics.* Pre-theoretic classical reasoning is used to reach conclusions about their observations. One problem is that observers are not truly justified in claiming that their different measuring results are related to the same system. They may both view the same outcome value on the screen but there is no guarantee that this wasn't caused by two different systems, they simply assume this to be the case. To be sure of this the observers would have to expand what they count as the measurement to probe deeper into the environment, doing so would leave the measurement irreversible and perturb the system (thereby failing the intersubjectivity condition). Another assumption needed to make quantum darwinism work is that the observers have to recognise each other as distinct observers and the system as being well defined. However, this essentially amounts to the emergence of classicality having already taken place. Fields writes that "[w]hat is critical to [quantum darwinism] as an explanation of the emergence of classicality from minimal quantum mechanics alone is that O_1 and O_2 have no prior shared classical criteria with which to identify S, since the availability of any shared classical criterion sufficient to identify S would imply that S had already 'emerged' into classical objectivity" (2011, p 4). These sort of objections are common to the decoherence program as a whole, not just to quantum darwinism; we have to define the split between the system and the observer to be able to make the derivation.[5]

However, these criticisms use a very specific kind of observer: an observer that is making distinctly quantum measurements. But decoherence, and quantum darwinism in particular, is not primarily an account of the emergence of quantum measurements. Rather it is about the emergence of a world that we are able to observe in an intuitive, everyday sense without having to explicitly account for quantum mechanics. An explanation of what a observer is and how they make quantum measurements is a secondary task that is not taken to be a core part of the explanation that decoherence provides. This is true of all the decoherence programs, for example in the decoherent histories program Gell-Mann and Hartle (1994) explore how IGUSs can be added to explain how observers evolve within a decoherent world (see also Dowker & Kent 1996; Rocha, Rickles, & Boge 2021 for discussion); but this is an add on with decoherence assumed to take place independently. Quantum darwinism, despite its framing in terms of observers, is no different. Durt (2010), along the same lines as adding IGUSs to decoherent histories, argues that observers who operate on entirely classical logic evolve by relying on the decoherent structure that quantum darwinism explains. The world emerges as classical without any reference to observers or measurements; only once it has emerged can observers operate.[6] In this sense Field's charge that observers assume that the system has already emerged is entirely correct. But it is not a problem for quantum darwinism as a theory of emergent classicality in general; the task of explaining observers is a secondary one and quantum darwinism should not be taken to stand or fall on this issue.[7]

The rest of this chapter will show how the records that quantum darwinism produces should be understood as analogous to the everyday record systems we see in the classical world and not as something closely connected to quantum measure-

---

[5] Relatedly it is frequently charged with assuming that the Born rule describes outcome probabilities (e.g. Okon & Sudarsky 2016).

[6] One thing to note is that what may be causing confusion here is what counts as an observer. A thin sense of observer is simply as an instrument that performs measurement. The implication here is that being an observer requires pre-measurement beliefs and a much fuller mental state - it is this sort of observer that we can view as belonging to the classical realm.

[7] The criticisms from Field *are* concerning when it comes to this secondary project, although the most recent presentation from Zurek 2022, while not addressing them explicitly, may show progress on this issue.

ments. These records are the sort of epistemic structure in the world that allows observers to operate on a day to day basis, but they are not dependent on or necessarily linked to specific measurements. Regardless of whether or not quantum darwinism can fully model quantum measurements, it *can* describe how this standard of classicality emerges. Explaining quantum measurements is undeniably part of the literature on quantum darwinism, but recognising that it is not essential to the project will help to direct criticisms in a more productive way.

5.2.3 Beyond Diagonalization of the Density Matrix

Before delving into how records act as a standard for classicality in more detail I will take a step back and compare the standard account of decoherence to quantum darwinism in a scattering example. What this will show is how records indicate robust classical dynamics in a similar way to the conditions of decoherent histories. Notably, this does not involve any actual measurements but just the conceptual possibility that something could be a record. It will also bring forwards a key feature of quantum darwinism: redundancy.

Photons scattered off a central atom in a position superposition decohere the atom and encode information about its decohered position. In the case of scattering with a single photon it is not clear how we should regard the final state of the atom. If the photon has the appropriate wavelength (roughly corresponding to the distance between the two states of the superposition) then the photon becomes strongly correlated to the atom and the reduced density matrix of the atom is effectively diagonalized. This is considered a case of decoherence in Wallace (2012) and Joos *et al* (2013). In both cases the single photon case is considered only briefly and then disregarded in favour of analysing in more depth the multiple photon case. Additionally, decoherent histories are introduced to more explicitly consider the emergent classical dynamics that result from multiple scattering events. After interacting with a single photon the atom *can* be shown to follow classical dynamics; but while this is technically possible the situation is very unstable. It would be relatively easy for the photon to re-interact and unentangle from the atom, reversing the process of decoherence. It would also be easy for the atom-photon correlation to be disturbed in some way and the dynamics of the reduced density matrix with the photon traced out would not be representative of the state of the atom. Any expected classical behaviour of the atom, if it arises at all, is likely to be short lived and unstable. This does not match with the robust and stable classical world we are familiar with.

Quantum darwinism approaches the problem differently and shows that asking whether or not the photon can act as a record leads to similar conclusions to decoherent histories; namely that multiple scattering events are needed before the atom can be considered effectively classical.

First let us look at the single photon case. Practically creating a device that can extract information from a single photon is not obviously possible. And certainly it would not be possible for multiple observers to independently measure the same photon. But even removing the technical considerations, we can focus just on the idea that the photon is, at least nominally, a record regardless of our ability to read it. And we can show that in the single photon case this is not true - thereby showing that according to the standards of quantum darwinism the atom is not decohered. Zurek's (1982) early work on decoherence proved this by looking at interacting a two state system - an atom with excited and ground states - with a system of interest - a spin system that could be up or down. The atom system acts as a rudimentary measuring apparatus; entangling the two systems should in principle allow us to extract information about the spin through its correlation to the atom state. However, Zurek shows that this is not possible and the atom cannot act as a record. This is because the states of the atom system and the spin system can both be independently rewritten in many other bases. For example a spin in the $|\uparrow\rangle$, $|\downarrow\rangle$ basis can be rewritten in the basis $|\rightarrow\rangle = (|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$, $|\leftarrow\rangle = (|\uparrow\rangle - |\downarrow\rangle)\sqrt{2}$. Similarly for the atom system, the basis of ground $|\underset{\circ}{=}\rangle$ and excited $|\overset{\circ}{=}\rangle$ states can be rewritten as $|+\rangle = (|\overset{\circ}{=}\rangle + |\underset{\circ}{=}\rangle)/\sqrt{2}$, $|-\rangle = (|\overset{\circ}{=}\rangle - |\underset{\circ}{=}\rangle)\sqrt{2}$. As such the state of the atom cannot be a record of the spin state; the atom state $|\underset{\circ}{=}\rangle$ could correspond to a spin state of either $|\uparrow\rangle$ or $|\rightarrow\rangle$ (or any other possible basis) and hence tells us nothing. For it to be a record there must be a correlation within a specific basis. This conclusion can be applied to the case of a photon scattered off a central atom: they are entangled but not in an informative way.

Now we turn to the multiple photon case. When multiple photons (which can either be energetic enough to each individually strongly entangle with the atom or less energetic such that individually they only weakly entangle with it, but collectively they decohere the system) are scattered off the atom, then a more complex array of measurements can be performed to locate that atom. Different observers can intercept different sets of photons to independently obtain their results. So the multiple photon case satisfies the information account of emergence. It also satisfies the dynamic account: the possibility of re-interaction (of *all* the photons) or disturbance is negligible when multiple photons are involved, and we can expect stable classical dynamics. Zurek (2022) shows that the redundancy of records is equivalent to what he calls *redundant decoherence*. This is where not just the reduced density matrix of the atom with respect to the environment is diagonalized but also the joint state of any atom-photon pair (tracing out the rest of the photon environment). This corresponds to a correlation in a specific basis (in other words the selection of a pointer basis) and means that the photon is now an informative record. It also means that the atom's density matrix remains consistently diagonalized, and follows classical dynamics robustly. Even if one photon re-interacts or is disturbed, the other photons are sufficient to ensure that the atom can still be accurately represented by the effectively diagonalized reduced density matrix and follows classical behaviour.

Quantum darwinism helps to identify cases where classical behaviour is particularly robust and stable beyond just the weak instantiation of classical behaviour that the effective diagonalization of the density matrix ensures. It provides a more stringent standard for when classicality has been achieved and emergence can be justified. It is doing the same job as decoherent histories but introducing a novel way to model this effect.[8] Records are used as a benchmark to measure against, to determine when the decoherent state is achieved; actual observers looking at these records are not necessary at all.

---

[8] Largely the work on quantum darwinism is separate from decoherent histories and makes little mention of it, except for the paper by Riedel, Zwolak, and Zurek (2017) connecting them. However, Gell-Mann and Hartle's work is often cited as one possible idea for solving this sort of problem.

5.3 What are records?

5.3.1 A General Account

I now turn to look at what a record is outside of the quantum context. Doing so will give a better handle on what classical features we are looking to derive via quantum darwinism. In the previous chapter I argued that records should be defined as correlations that are robust against noise; this robustness is an essential property of records and not just an instrumental concern about how they are used.[9] Here noise is defined using ideas from signalling theory: it is irrelevant modifications to the record that disrupt or obscure the information it encodes. This follows from considering how records can work in non-idealised settings and how we are practically able to identify informative correlations in the world. With this in mind it becomes clear that robust, redundant correlations are a common feature of the classical world; this feature is what allows us to make everyday inferences about ordinary objects. Quantum darwinism should be understood as showing how these sorts of correlations emerge.

Robustness is achieved by two factors: 1) the use of stable, macroscopic systems with reliable evolutions, 2) the use of redundancy. The first condition largely acts in the background to identify the sort of systems that can be robust in the necessary ways. It also immediately suggests a connection with emergence. But it is redundancy that I will largely focus on here as this is where the link to quantum darwinism will become clear. Redundancy works because for a single copy of the record we cannot identify whether it has been affected by noise or not. Creating multiple redundant micro records means that although each individual micro record cannot properly be called a record as it does not give reliable information, taken together the redundancy of the information ensures that a robust correlation holds.[10]

---

[9] When thinking about records and time we often explicitly talk about how they allow us to make inferences about past events rather than just about other systems. This was the focus in the previous chapter. However, lasting over time is just one aspect of records and records that are correlated to the current state follow the same set up.

[10] The word redundancy is often taken to mean eliminable and playing no essential role. The argument here is that the many copies *are* essential. They are redundant in the sense than any single copy is eliminable provided the overall redundant set remains.

Given its importance to the argument of this chapter, I will spell out how redundancy works in everyday examples of records in a little more detail. Consider a footprint on a muddy path. Mud is rarely particularly smooth. It is normally created by, for example, the ground being disturbed by many people or animals passing through. An already churned up muddy path is a difficult record to read. It contains many 'records'; but most will be unidentifiable as anything specific. If an identifiable footprint can be found, then it is because we see a shape that has a clear correlation to the shape of a boot that stands out from the general noise of dips and troughs around it. This robustness is important for identifying the record. A partial footprint may still be identifiable but the less of it that remains the harder this becomes. The different parts of the footprint outline form a sort of redundancy.

There are two ways redundancy is working here. First there is a straightforward sense: only a partial footprint is needed for us to be able to identify it. The imprint of the heel, maybe with the pattern from the sole of the shoe, is often enough to make it clear that this is a footprint. The entire footprint therefore contains a redundancy of partial footprints and this redundancy allows for robustness against noise. Even when partially destroyed the record remains informative. One might object to using the term redundancy in this situation. It is not straightforwardly many identical copies of the same information; instead many similar correlations are used. However, it is still the case that we have a surplus of information. Each fragment is possibly informative, just as a single micro memory device in quantum or nano-computing can encode information. But it could easily be obscured by noise. Adding many similar records verifies the information in each one and reduces the chances of errors.[11] The second sense of redundancy is less obvious and is required for the individual parts to be informative at all. Each fraction of the outline is correlated to the foot and contains some information about it in virtue of this. But there is a minimal size of fragment that is needed for us to identify the mark as being a footprint at all. The required size is somewhat context dependent: an expert tracker needs a smaller fragment than a layperson and some parts of

---

[11] Further redundancy can be found when, for example, the footprint is part of a set of footprints one after another.

the footprint are more recognisable than others. But below a certain size it is unidentifiable; it is just a random mark in the mud. The fraction itself remains unchanged and it is still correctly correlated to the foot, but no information can be extracted from it. The correlation cannot be separated out from noise. Only by collecting together many similar fractions - a redundancy - to form a larger fragment can a record be formed. The redundancy of insignificant correlations creates a robustness that stands out from the background noise. While a small mark may be the result of anything a distinct line is much more clearly the result of one thing or another. Another example of this sort of redundancy is pixels in a photograph. Each pixel is correlated to the subject of the photograph but a single pixel can tell us little to nothing. It contains only very minimal information and could easily be affected by noise. But the pixels taken together create a highly informative photographic record. The information that a certain pixel is blue is verified by its placement in a whole set of blue pixels that show us a picture of the sea. However, a single black pixel among these blue ones tells us there was probably a speck of dust on the camera.

A certain amount of redundancy is needed for any information to be extracted, and once this minimal level is reached further redundancy helps to protect against errors. All records seem to have this feature. Sometimes it varies as to whether we label the entire redundant set as one record (for example we might say that a footprint is a single record or a photograph is a single record) or whether we count all the many parts individually as records (each pixel in the photograph is a record). This difference is arbitrary and redundancy features regardless. In the former we say the record is a record because it has many redundant parts. In the latter we say the record is a record because it is part of a redundant set; when taken independently of that set it is not informative.

5.3.2 Records as a Standard for Classicality

How does this account relate to quantum darwinism? The initial starting point for applying the general account of records to quantum darwinism is that there is a striking similarity between Zurek's language of redundancy and robustness and the way that records operate in a classical world. The next section will explore this in detail and carefully consider whether the account given above accurately matches the records in quantum darwinism. Before I look at this comparison in detail, however, we can first consider how classical records act as the standard of classicality that quantum darwinism aims to derive.

The dynamic account of decoherence defines the classical world as a world that follows classical dynamics. Quantum darwinism defines it as a world in which observers have intersubjectivity and can make repeat measurements without significantly affecting the system. We can treat ourselves as separate from the world around us because we can observe it without influencing it; two people can look at the same object and agree what it is. In a more scientific context, multiple experimentalists can read off the value on the screen. This is the feature that quantum darwinism aims to explain, and it follows directly from the *robustness* of the epistemic structure in the world, a feature which the account of the previous chapter foregrounds. Robust records are the basis on which we are able to make regular inferences about systems without having to take into account a formal theory of measurement. Observers in the everyday world never have to think in terms of quantum mechanics or model themselves explicitly as a system taking part in a measurement; any interaction that we have with the system is counted as irrelevant noise. The irrelevance of the observer to the state being measured is so ingrained that we rarely even explicitly think about it.[12] The negligible impact that observers have on the system is also what makes intersubjectivity possible. The first observer leaves the record unchanged for the second observer to look at. The robustness of records guarantees that any effects that come from interacting with the observer will be irrelevant in this way.

Records are found naturally all around us in all sorts of systems, frequently appearing outside of explicit measurement contexts. We can and often should model records like a result displayed by a measuring apparatus in this way, but records are not limited to this. They are a generic and frequently found feature of the classical world. Criticisms of quantum darwinism, such as those by Field laid out above, define records as the result of measurement and model this in the way commonly

---

[12] Of course there are classical cases where our measurements affect the system; but these are outliers and do not disprove the general assumptions that we make.

done in quantum mechanics: a measuring device is written as the state $|ready\rangle$ before it interacts with a system and is transformed into a record state such as $|up\rangle$. The way that records are used in the decoherent histories program is also very close to this. The records are global projection operators that allow one to distinguish between different histories - where the histories are sequences of distinct measurements. However, this is directly analogous to Albert's ready state account of records that I critiqued and rejected in the previous chapter. It does not provide a realistic model of the features that records must have to count as informative. When we model a measurement in this way we ignore all the important factors that go into making a measuring device robust and reliable.

Quantum darwinism should be understood as deriving how robust correlations acting as records form in quantum mechanics; not as a result of an specific measurement from an observer but just because of the general interaction with the environment. The presence of observers explicitly studying quantum mechanics in that environment is happenstance and the process happens with or without them. It is the formation of these robust correlations that allows such observers to operate in the way that they do. Even when a measurement *is* the cause of the decoherence this still does not prohibit us from understanding the physical measuring apparatus as just a type of environment. Despite our tendency to describe the device as a single quantum state $|ready\rangle$, we should really consider all the parts as a compound environment that generically have this sort of effect on a system.[13] What makes it a measuring device is the final robust correlation between a record displayed on the screen and the state being measured. It is this that makes it informative; but the only difference between this and a naturally occurring environment is whether or not we intended in advance for this to be a measurement. The actual process is the same. The robust correlation is the result of all the redundancy within the system.[14] At this point it is also obvious that the record can be read out many times

---

[13] However, it is true that the current practical examples of environments studied in the quantum darwinism literature tend to be simplistic environments made of a set of dynamically independent, identical subsystems. The sort of environment represented by an actual measuring device would be far more complex and it has not been shown explicitly that the same results would hold. For now we can only extrapolate that this general process will still apply.

[14] When we have multiple observers with multiple measuring devices this still holds. Fields (2010) raises the objection that multiple observers have to assume their results are correlated to the same system and that this genuinely is redundancy; this is certainly true and the experimental

by different observers. The photons carrying the information between the screen and the observer *technically* dynamically entangle the systems but the influence is so negligible that it is dismissed as noise and we treat the interaction as classical; it has no effect on the record displayed on the screen.

Understood in this way, quantum darwinism delivers an account of how a regular and essential feature of the classical world - robust, redundant, informative correlations - emerge. Observers are secondary to this and rely on this sort of structure already existing in the world. In this light the criticisms levelled against quantum darwinism, which claim it presupposes classical observers and well-defined systems, are far less concerning.

5.4 Records in Quantum Darwinism

It still remains to show, however, that the general account of records I have given is actually an accurate description of what is going on in quantum darwinism. There is a strong initial similarity: the creation of redundancy in quantum darwinism is what singles out a single correlation (the pointer basis) and makes the individual records informative. As the photon scattering case study showed, an individual photon is not a record unless it is part of a redundant set all decohering the same central atom. This use of redundancy seems the same as the redundancy described in records generally. But what remains to be shown is that this can be connected to the definition of robustness against noise. Noise has not so far featured in the analysis of decoherence presented here, so this similarity of using redundancy to identify a correlation needs more careful examination before a comparison can truly be drawn. The following subsections, therefore, will lay out exactly how records - defined as they are above as correlations that are robust against noise - arise out of quantum darwinism. Showing this will also serve a dual purpose by setting up a lot of the arguments needed to show that records are related to the screening off

---

methods take careful designing to isolate the desired target system. However, the dynamical process that creates the redundancy necessary for measuring devices to function justifies this. Zurek (2022) also gives a much more detailed and up to date technical presentation of how quantum darwinism actually treats observers. This involves multiple stages, and a number of different factors contribute to what we think of as an observation - the observer is never modelled as a simple quantum state.

condition of emergence that section 5.5 will draw together in more detail.

5.4.1 Redundancy and Noise

Care must be taken when connecting decoherence to noise. In the initial quantum superposition of the system (prior to decoherence through interaction with the environment) there are quantum correlations that are phase relations between subsystems of the Hilbert space of the system. These manifest as interference effects. Joos (2000) rightly cautions against any interpretation of these correlations as noise and of the superposition as an ensemble. Doing so would imply that the system is in a definite but unknown state and the quantum correlations are simply obscuring it. But this would not account for the presence of non-local features of superpositions, specifically the presence of interference effects. Noise obscuring an unknown state would not result in these effects.

Only after decoherence has occurred should we consider the analogy with noise.[15] The insistence from Joos (and also Schlosshauer 2019) that decoherence should not be considered in terms of noise largely stems from attempts early in the literature to separate decoherence out from other processes in quantum systems, in particular from dissipation. Dissipation - energy loss from noise processes such as fluctuations and imperfections in the system - can have a similar effect to decoherence in damping the off-diagonal terms of the density matrix. In contrast, decoherence specifically describes the effects of entangling the system with the environment and the loss of information this produces. During the process of decoherence the system becomes correlated (entangled) with the environment and some of the quantum correlations in the system are replaced by the environment-system correlations - this process singles out a pointer basis. The original phase relations from the system superposition are delocalized throughout the total combined system-environment (and eventually throughout the entire universe) (for more detail see Zurek 1982; Joos 2000; Joos & Zeh 1985). The information that quantum correlations hold is

---

[15] This should not imply a troubling circularity where records are used to define decoherence and then decoherence is used to classify the resulting states as records. Decoherence produces certain states through the dynamical interaction and these states then become classified as records. This is the same in dynamic decoherence in which classical dynamics are derived and then this is used to classify the off-diagonal terms in the density matrix as irrelevant.

spread throughout the entire combined system, while records of the pointer observable are created in many parts of the environment. This corresponds to the off-diagonal terms in the density matrix of the system-apparatus (with the environment traced out) tending towards 0. The state can then be treated as if it was an ensemble of measurement results (although the measurement problem - how a single result gets selected - remains unsolved; but this is not relevant to the problem at hand here). Decoherence can be shown to proceed on a different timescale to dissipation.

The association between noise and dissipation comes from the specific usage of noise that is common in that literature, mainly referring to fluctuations. However, the definition of noise given in section 5.3 is more general and can encompass a wider range of effects. Noise is defined not just as fluctuations and random effects but more generally as irrelevant contributions to the record that obscure the robust correlation between record and system. This encompasses a wider range of phenomena including systematic low level details that have negligible effect on the overall state of the system and its evolution. Also notably when applying the concept of noise to records this is *after* the environment has become correlated to the system. The process does not *proceed* through the effects of noise. But this does not mean that we cannot apply the generalised concept of noise to the quantum correlations that are delocalized throughout the environment after decoherence.

Despite it being important to distinguish between dissipation and decoherence, however, dissipation also contributes to emergent classicality. The decoherent histories frameworks explicitly considers how *all* sorts of effects contribute to deriving classical dynamics. While decoherence was originally focused primarily on quantum correlations and the effective diagonalization of the reduced density matrix, it now extends beyond that. The combined effects of decoherence, dissipation, noise, fluctuations etc are all included. The records that quantum darwinism produces naturally subsume all of these other influences under its standard of emergent classicality. Using *noise* as a generic term that captures all such effects makes this very clear. This is another way in which quantum darwinism is doing a very similar job to decoherent histories but solving the problem in a different way.

However, there is still the challenge of showing that the quantum correlations (interference terms) specifically are irrelevant and can therefore be classified as noise. This strongly connects the project to accounts of emergence in the philosophical literature. The irrelevance of lower level details is a commonly cited criterion for dynamic emergence. Wallace (2012) and Franklin (2022) both show that the interference terms are irrelevant for the evolution of the resulting classical state. I will not repeat their full analysis here, although it will be discussed further in section 5.5. What they show is that the system - described by the reduced density matrix with the environment traced out - is effectively diagonalized and obeys classical dynamics and probability calculus. This justifies treating the interference terms as irrelevant.

Much of their analysis is directly applicable here. However, there are two issues with carrying it over to justify the irrelevance of the quantum correlations to records. First, quantum darwinism shifts the attention to the environment and no longer traces it out to leave the diagonalized reduced density matrix. The justification for considering the interference terms as irrelevant focused on the system itself; this justification does not necessarily carry over to the environment. Second, the irrelevance to the dynamics does not automatically help us understand the information-theoretic context of records in which the dynamics do not play any major role. Instead it must be shown that the quantum correlations are noise with respect to the correlations in the environment that observers measure.

5.4.2 Local versus Global

We can find our justification by considering the difference between the global and the local environments and how only the latter is relevant to records. When taking the environment as a whole the interference terms delocalized throughout it cannot be ignored or dismissed. Various authors stress this point. Despite the irrelevance of the interference terms to the reduced density matrix of the system, the interference terms are still relevant for the joint state of the system-environment. The total joint state remains in a pure superposition. Joos & Zeh (1985) state that "The interference terms still exist, but they are not *there*" (p. 224). By this they mean that the quantum correlations (entanglement relations) found in a superposition state (i.e. the interference terms) are not destroyed when a system is entangled with the environment. Instead the correlations are spread throughout the entire system-apparatus-environment combined system. While the reduced density matrix of the system is effectively diagonalized, the overall state is not; the interference terms are still present even if they are not in the original system. As such, these terms are still highly relevant for the dynamics of the total system. In particular the entire combined state still obeys unitary dynamics and is fully reversible. An operation over the whole system and environment would allow an observer - at least in theory - to reverse decoherence. Given that the full state of the system-environment retains the superposition, Joos (2000) stresses that the interference terms in the reduced system cannot be regarded as irrelevant noise. Additionally, moving away from dynamics and back to the information theoretic framework, a measurement of the whole environment and not just a fraction of it would allow an external observer to observe interference effects as well as gain information about more than just the pointer observable. In the context of quantum darwinism such a global measurement would fail the classical standard as it would disturb the superposition and would not be verifiable by others.[16]

To this last point we can reply with a practical note: such a global measurement may be definable in principle but practically would be very difficult to achieve. It would require the large environment to be perfectly isolated from the observer else the observer would be part of the environment and have to include themselves in the measurement. In many cases we consider the spread of decoherence through the environment to eventually spread throughout the whole universe. Any realistic observer within the universe would be unable to perform measurements at this scale. Even considering a smaller, more limited environment where a global measurement is more feasible, quantum darwinism requires that the information is redundant enough for observers to potentially perform individual measurements on different fragments of the environment. Therefore, a global measurement on the whole environment would have to be on a significantly larger scale than any

---

[16] A measurement could be seen as the observer simply becoming part of the environment itself and entering into the entanglement relations. The point remains the same that a second external observer could not now just measure the environment and corroborate the first observer's results. They would also have to include the first observer in their measurement as part of the environment.

common place measurement on just a fragment. The possibility in principle of the global measurement does not relate to the practical considerations of how we make observations in the world. The latter was the motivation for quantum darwinism; it aims to explain why we make observations of a classical world all around us. Commonplace observations are made locally on fragments of the environment.[17]

Moreover, despite being *in* the environment, records do not make up the whole of it. In fact, as emphasised by the introduction of redundancy, only a small fragment of the environment forms a record and this is repeated many times. Meanwhile the relevance of the quantum correlations in the environment - both to the dynamics and to global measurements - requires the *global* combined system. The essence of quantum darwinism is showing that measuring fragments of the environment puts a classical limit on the information that can be extracted (only information about an observable is available). The global superposition state is irrelevant to the local fragment.

The local scale is definitional of classicality. Looking at the notion of classical information, what distinguishes it from quantum information is how parts relate to the whole. In the classical realm knowing a composite system implies knowledge of its parts. This does not hold true for quantum systems where the state of the whole cannot tell us about the parts; the parts do not necessarily have well defined states at all. Measurements on a joint quantum state reveal more information than any sequence of measurements on its component parts. But even for a quantum system, local measurements allow us to extract a clearly defined subset of the total information: the classical subset. This is information about observables rather than the state itself. For each individual record making up part of the environment the information it gives is purely about the classical state of the system.

On top of this Zwolak and Zurek (2017) show that many parts of an environment are irrelevant to the redundancy of the correlations that form the records. They

---

[17] This may be good reason not to consider the records used in decoherent histories to be genuine records. As Gell-Mann and Hartle (1995) note, these generalised records are not actually accessible to observers and are global projectors. We could not consider the quantum correlations as irrelevant to these ‘records’ as they use the global state of the system and environment. In this way quantum darwinism gives a much more detailed account of how actual records are produced and play a role in decoherence.

are considering a complex environment where many parts of it do *not* become correlated to the system at all (and carry no information about it). They show that the existence of these parts of the environment does not have any appreciable effect on the redundancy of records in the environment and how likely an observer is to intercept at least one. As such, they argue these parts of the environment are irrelevant to quantum darwinism.

Based on this, the quantum correlations spread throughout the environment can be regarded as noise - in the general sense - with respect to the local records that an observer would measure. The conceptual possibility of a global measurement or operation is not relevant to these local records. The quantum correlations are irrelevant contributions to the final state of the record that is determined by its robust correlation to the system.

This provides justification for claiming that the redundancy of records in quantum darwinism protects the system against noise. This matches both senses of redundancy that were described in the context of the footprint example at the end of section 5.3. First, redundancy has the effect of making a particular correlation identifiable over and above the many other correlations; this is the active process where the creation of records suppresses the other possible bases in favour of the pointer basis. Due to this process the correlation between the observable and the record becomes identifiable against the background of many quantum correlations delocalized throughout the global system-environment state. The latter can therefore be characterised as noise with respect to the former. This is more active than the classical case as the record creation has an appreciable transformative effect on the system while in the classical case it is often passive (the mud does not transform the shoe stepping in it - in most cases at least). But the overall role of redundancy plays the same role of making information identifiable. Second, further redundancy allows for robustness against disturbance by noise in a much more straightforward sense. Here noise is the negligible contributions to the behaviour of the system made by the delocalized quantum correlations, but it can also be the effects of observers interacting with the records and making measurements, and finally it can also be the more specific meaning of noise as fluctuations. An observer can robustly and consistently extract the classical subset of information about the state from the records and ignore the rest of the environment (and the quantum correlations spread through it) as irrelevant to this information.

5.5 Records and Emergence

What is left to explore is how records as a standard of classicality relate to the philosophical work of pinning down what emergence really means. This ties together various threads that have already arisen in the previous section. What will be shown is that records can be used to justify the conditions of emergence as an alternative to using the dynamics of the system. In any physical situation the two will go hand in hand (after all, records are created through a dynamical interaction); there will be both the production of records as described by quantum darwinism and the emergence of classical dynamics. But what this shows is that information provides a valuable alternative route to understanding emergence.

Philosophical accounts of emergence in science largely focus on the autonomy of the higher level theory and how it screens off irrelevant details of the lower level theory to provide novel explanatory or predictive power (Woodward 2021; Ross 2000; Ladyman & Ross 2007; Franklin & Robertson 2022; Robertson 2021). The emphasis is on showing how emergent theories are objective and independent levels of reality, not just pragmatic approximations of fundamental physics.[18] In other words it answers the question of why we should use any other theories except fundamental ones.

Emergence has been defined in terms of *conditional irrelevance* (Woodward 2021; Franklin & Robertson 2022). This is not specific to quantum decoherence but applies generally to the emergence of higher level theories and ontologies.[19] Conditional irrelevance explains why it is possible to screen off lower level details to focus on the instantiation of the higher level theory. To understand it we can first consider its counterpart: unconditional relevance. When accounting for the

---

[18] Accounts of emergence developed in the philosophy of physics literature differ from how emergence is considered in other areas such as philosophy of mind. Emergence is considered compatible with reduction and the mathematical and conceptual compatibility of the two theories remains important.

[19] How exactly to formulate the conditions of emergence is contentious; the view described here is what I take to be the most prominent account.

behaviour of a system, it is clear that the microscopic details that the system is made up of will be relevant to the explanation. As the system evolves in time it successively instantiates a series of microscopic states. Hence these microscopic states are unconditionally relevant to explaining the system's evolution. However, this description may be very unwieldy, long winded, or complex. Due to measurement limitations the microstate is likely to be impossible to determine and even if we could, the resulting calculation to get from the initial to final state may be practically impossible. This is not what scientific practice attempts to do. Instead we note that many of these microscopic details do not have any significant effect on the overall evolution of the system. We can more effectively describe the system by conditionalizing on certain higher level features or properties. The behaviour of the system can now be explained more simply and efficiently by using them. These features could be instantiated by many different possible microstates; hence the lower level details are being screened off and become *irrelevant* to the dynamics of the system. But only once the higher level features are accounted for, hence *conditionally.*

Franklin (2022) applies this account to the problem of emergent branches (assuming an Everettian interpretation) in quantum mechanics. He argues that the off-diagonal interference terms in a decohered system are conditionally irrelevant to the dynamics of the reduced density matrix.

> "decoherence underwrites the screening off of otherwise relevant details and is, thus, responsible for the screened off states' evolving according to a novel (quasi-)classical dynamics." (Franklin 2022, p. 7)

Screening off, in the form of conditional irrelevance, is essential to emergence. The dynamic account uses dynamical laws to justify it. The reason dynamics are commonly picked out for emergence is that they provide a way to make concrete predictions over time that can be tested against the actual behaviour of the world.[20]

---

[20] For example, Franklin (2022) looks at a case study of the dynamics of Hyperion and compares the computer generated dynamical models in phase space to show that the decohered dynamics are the same as the classically modelled dynamics. There are, of course, theories of emergence that do not focus on emergent dynamics (or example, the emergence of phase transitions in thermodynamics). However, these tend to focus on specific entities instead. Quantum darwinism seems to do neither.

Simply put, the question of emergence is ‘why does a fundamentally quantum world *appear* classical?’; if we can derive classical dynamics and show the world behaves as it predicts, then we have an answer. From this follows the claim that the world is genuinely classical at a certain level of reality.

However, records can equally indicate that screening off occurs. By the definition of records put forward they must be robust against various noise effects. These include fluctuations and the effects of multiple interactions. It also includes general low level details that are irrelevant to the dynamics. The definition of noise proposed is in exactly these terms: noise is any irrelevant modifications to the record. There need not be fear of circularity here of having defined records in precisely the terms needed for this argument to follow. This definition of noise, and the importance of noise to records, followed from much more general considerations about what a record is and drew on a definition of noise common in signalling theory.

Records must have this feature because realistic records must be accessible to observers (even if no observers are actually involved this is part of how we can get a handle on what records are). Accessibility requires them to be local - as discussed in section 5.4.2; this makes some kind of screening off necessary. The global state must be irrelevant to the local variables that form the record (and its correlation to the past state it is recording) so that the local record is able to give information. Conditioning on the information in the record, all these details are irrelevant to determining the state of the system.[21] The previous chapter discussed how using records to find out about the past or the state of another system is distinctly different from using retrodiction. Retrodiction gives information by taking the entire state of the system and evolving it backwards under the dynamics. In practice of course this is rarely done in completeness. Instead shortcuts are taken to make the calculations possible, such as making models of particular phenomena using idealisations and other approximations, all of which allow us to ignore various irrelevant aspects of the world when making retrodictions and predictions about the phenomena. Much of the literature on emergence is dedicated to showing that these

---

[21] We can also see that the connection between records and screening off depends on the selection of specific variables of interest as well as spatial locality. For example the cosmic microwave background radiation spans the entire universe but the record itself is only the thinly spread photons from the big bang and we screen off all the other objects that are emitting photons.

sort of methods lead to emergent theories. Records essentially take the use of such shortcuts to the extreme. They require that a specific feature of a local system is able to give information; all the details of the system outside of this specific feature can be ignored. External influences need not be taken into account and can be treated as noise when making inferences about how the record is correlated to the aspect of the system it is recording. In essence, these details are screened-off. This is what makes a record a record and separates it from the general process of retrodiction. This locality leads to the definition of records as correlations that are robust against noise.

The work done in section 5.4 has shown that we can treat the wider state of the system-environment as irrelevant to the the state of the records of the pointer observable that are created in many fragments of the environment. This is directly analogous to the dynamic account of emergence from Franklin (2022) that shows that these details are irrelevant to the dynamics. The screening off condition underwrites both the dynamical accounts of emergent classicality and the records that define the information theoretic standard of classicality.

5.6 Conclusion

To summarise I return to the three aims laid out at the beginning of this chapter: 1) to further explore the important role that records play in physics and to showcase how the account I developed in the previous chapter can be used to strength these theories and improve our understanding of them, 2) to prove that information can be used to understand emergence, and 3) to show how quantum mechanics fits into the model of time I am developing.

For the first aim: I have argued here that records as robust correlations are the standard of classicality that quantum darwinism aims to derive. When understood in this way it becomes far clearer what quantum darwinism achieves and what exactly it means for classicality to emerge in this context. It also suggests that the criticisms levelled against quantum darwinism regarding its use of observers need not be a concern for the project as a whole, only for the secondary task of using quantum darwinism to model quantum measurement outcomes. Although this is not a complete defence of quantum darwinism it provides a stronger and clearer understanding of it. This clearly illustrates how important epistemic structure such as records can be to the way that we use and interpret physics (and the specific utility of the account I have developed), hence providing strong motivation to include them in our metaphysical thinking as well.

For the second aim: I have shown how records can be connected to philosophical accounts of emergence through the screening off condition. This is commonly justified using the dynamics of a theory, but records can be an alternative way to identify this. The role that records play in quantum decoherence is evidence of how central epistemic structure is to our theories, and how closely this is connected to emergence. What can be known and how things can be measured is a defining feature of different levels of reality; we cannot assume that these notions are basic or can be presumed at any level. Looking at how epistemic structure emerges gives us a much better understanding of intertheoretic relations. Additionally the discussion of how records are local helps connect emergence to the ideas about perspectives used in this thesis. Considering how localisation leads to the conditions for emergence shows how perspectives, in the sense of being localised to a particular position, are connected to perspectives in the sense that we are at a macroscopic level of reality.

When thinking about time this can help us make sense of what it means to say that the flow of time is emergent. The literature on emergence gives clear criteria for how to think about the relation between metaphysical claims of theories at different levels. Connecting time to this literature, through the use of records, allows us to understand how flow can emerge even if there is no metaphysical difference between past and future at the most fundamental level.

Turning to the third aim, I will finish by addressing how the analysis of quantum darwinism fits into and informs the model of time developed here. What has already been said regarding the first two aims shows how it provides clarity on the question of emergence. But there are more specific results from the literature on decoherence which strongly resemble the model of time I propose.

Decoherence is also often credited with the appearance of asymmetry in quantum mechanics and here again we see records playing a role in mediating between an open future and a fixed past. The work on decoherent histories has directly addressed how the presence of generalised records is essential, firstly for singling out what possible histories give a sensible quasi-classical macrohistory, and secondly for ensuring the "permanence of the past" (Gell-Mann & Hartle 1995). To expand on the latter: it is possible in the decoherent histories framework that future choices of coarse-graining can define different past histories, overwriting the history that is defined at the present time. But the condition of strong decoherence (the presence of global record operators that distinguish the set of histories) prevents this; this also ensures that histories defined relative to the present can be extended into the future. Quantum darwinism adopts the same framework (as argued in Riedel, Zurek & Zwolak 2016) but allows us to take the extra step in understanding that the past is fixed on the basis of real, local, records; this is no longer reliant on the much more abstract notion of a set of global 'record' operators. Additionally, on a more local level, it explains the repeatability of measurements by ensuring that a second look at the measurement apparatus or even a distinct second measurement will reveal the same result. This ensures intersubjectivity so that multiple observers can come together and find that they share a common, objective, past.

More generally, decoherence shows that quantum mechanics can be used to explain the macroscopic classical world that we experience, which has a fixed past and an open future. While I have focused on how decoherence fixes the past here, it is also used to derive the Born Rule, which justifies treating measurement outcomes as probabilistic (at least in interpretations that reject collapse and follow unitary evolution such as the Everettian or Many Worlds approach). Quantum darwinism in particular provides support for how we can mediate between ontic and epistemic readings of the wavefunction, an issue that is central to assessing the nature of quantum probabilities (e.g. see Wallace 2012, Franklin 2022). The state of a system - for example the reduced density matrix - encodes what properties the system can be said to have. But this should be seen as being partially defined by the correlations the system has entered into and how these correlations allow for epistemic access. Zurek says that:

> "one might regard states as purely epistemic...or attribute to them "ex-

> istence". Technical results described above suggest that truth lies somewhere between these two extremes, and these two aspects are complementary" (2022, pg 92).

Actively incorporating epistemic structure - by paying attention to the physical correlations in the world and how they can be used - into metaphysics once again gives us a better understanding of the way that physics uses probabilities.[22] We are no longer faced with the challenge of deriving the Born Rule - a rule of epistemic prediction - from an entirely ontic state. (This is something the Many Worlds approach struggles with and it has been a topic of much debate.) Instead we see that ontic states have an epistemic character all along; epistemic probabilities are indicative of physical structure.

---

[22] Zurek in fact goes further and develops this into a novel interpretation. Leading on from this Zurek reconceptualises objective existence to literally refer not just to individual objects but to objects in conjunction with a "halo" of information existing in the environment around them. Something cannot be said to exist without understanding how it is connected via robust correlations to everything around it. In the absence of such connections what status of existence we can assign to the state at all is questionable. Zurek calls these new objects, consisting of a state and its halo of records, *extantons.* This has the consequence that we should hesitate to attribute existence to the microscopic, non-decoherent, quantum state. He is also of the opinion that this implies that the universal state vector cannot be said to exist as it has no surrounding environment (and hence no halo). Robustness, which preserves a state through different interactions, is a necessary condition for existence. I will not go into extensive detail here and I do not outright advocate for adopting this approach in its full detail. It can be taken to give an entirely unique interpretation, as Zurek does (which he calls the existential interpretation); or it can just be seen as general evidence that ontic and epistemic go hand in hand and that our metaphysics should not ignore the epistemic relations in the world.

6 The Propagation of Definite Values in Relational Quantum Mechanics

6.1 Introduction

Localised temporal becoming is the moment to moment occurrence of events that constitutes the flow of time. It presents a picture of time that is dynamic and full of rich resources to explain our temporal experiences. Adding to this account, I have explored in previous chapters the role that records play in defining the difference between an open, indeterminate future and a fixed, determinate past. This view suggests a strong relationship between time and the accumulation of information. In this chapter, I will look at Relational Quantum Mechanics (RQM) - an information based reconstruction of quantum mechanics - as a particular case study to examine how localised temporal becoming can help explain how past values relate to present and future ones within this theory.

RQM is emerging as a prominent reconstruction of quantum mechanics that brings information to the foreground as the physical basis for the theory. It illustrates that rather than being an unimportant, merely derivative part of physics, epistemic ideas - and also the use of perspectives in connection with this - can play a central role in our interpretation of a theory and the way that physics characterises the past and future. However, RQM is still a relatively underexplored theory with many remaining conceptual problems, largely relating to how systems share information over time. Temporal becoming and its connection to the use of records may be able to help to make some sense of these problems. I intend to use RQM to show one way in which temporal becoming, and particularly the use of records in combination with it, can help us to understand physics. RQM provides a case study where this model of time is both useful and convincing, and working this out in detail articulates what the metaphysical picture looks like when applied to a particular theory.

However, I do not intend to tie localised temporal becoming to RQM alone. Temporal becoming can be used more broadly than that and may well fit with other interpretations of quantum mechanics as well. This chapter is meant to demonstrate just one way in which we can understand quantum mechanics in terms of information; and show how this supports the claim that the flow of time can be found by looking at epistemic structure. This builds on the work of the previous chapter on quantum darwinism, but shows that the information approach has multiple different applications. One need not accept either quantum darwinism or RQM specifically to recognise the importance of information and to see that there is a strong connection between time and epistemic structure that is found in multiple areas of physics.

I will first lay out RQM and the ontology behind it. Then I will present in detail the account of temporal becoming I will be using and show it works with RQM. Finally, I will focus on the role of records in creating a definite past and how this suggests solutions for some issues in RQM. These problems appear when RQM deals with sequences of interactions over time; in its current form it cannot adequately explain first, how predictions work for microscopic systems without a preferred basis, and second, how different systems communicate their previous results. I will discuss the recently proposed solution to this from Adlam and Rovelli (2023) and its shortcomings. As an alternative I will suggest that focusing on decoherent records would significantly restrict this problem. Ultimately, however, a full solution for the problem will rely on working out the ontology of RQM in more detail; without this RQM cannot make sense of how to model decoherence.

6.2 Relational Quantum Mechanics

6.2.1 States and Values

Relational Quantum Mechanics, originating with Rovelli (1996), is based on rejecting the ontological importance of the wavefunction. Instead, Rovelli takes inspiration from the original mathematical theory as it was formulated by Heisenberg (1925; also Born & Jordan, 1925; Born, Jordan, & Heisenberg, 1926) and independently by Dirac (1956). Heisenberg's work focused only on describing observable outcomes and made no claims about objects and their properties beyond this. This is the starting point for RQM. RQM takes definite outcomes to be the sole fundamental elements of reality and claims that there is no reality beyond this. This is combined with the postulate that there is a definite outcome at any interaction between two systems (including microscopic systems). The wave-function, which gives the continuous evolution of properties such as position (or the probability distribution of position at least), has no meaning beyond being a tool to store information about past interactions; information that is then used to calculate the probabilities for what values will be manifested in future measurements. The mysterious 'wave-function collapse' is an illusion; all that is happening is the updating of the predictive function with new information.

The concept of information is central to RQM and much of the interpretation is formulated in terms of information theory.[1] Primarily information is meant in the information-theoretic sense of correlation. If two systems are correlated then the total possible number of states for the combined system is less than the product of possible states for the systems taken individually. Applied to RQM: when two systems interact they exchange information (become correlated) and the value of a variable of one becomes dependent on the value of a variable of the other. As a result, those variables each become actual or definite relative to the other system. This combines the information theoretic sense of information as correlation with the slightly different sense of information as outcomes. RQM states that any correlation corresponds to a definite outcome regardless of whether a formal measurement has taken place. What makes it relational is that there is a definite outcome only relative to the systems involved and not to any third-party systems.

To explain this more clearly we can look at what happens during an interaction. Suppose we start at t_1 with a system s_1 in a superposition.

$$|\psi\rangle_1 = (\alpha\,|\uparrow\rangle + \beta\,|\downarrow\rangle)_1 \tag{2}$$

When a system s_2 interacts with s_1, s_2 finds a definite result of spin up or spin down (the wavefunction 'collapses'). Hence the state of s_1 *with respect to* s_2 (indicated by a '/' in the subscript) is either $|\psi\rangle_{1/2} = |\uparrow\rangle_1$ or $|\psi\rangle_{1/2} = |\downarrow\rangle_1$. The theory is fundamentally probabilistic, and these results will occur with probabilities α^2 and

---

[1] See Rovelli (1996) for details about the principles of information that RQM is built on.

β^2 respectively.

Now consider a system s_3 that does not interact with systems s_1 or s_2 during the measurement process. System s_3 possesses knowledge (through a prior interaction) of the initial states of the two systems. Using this knowledge s_3 can predict what will occur during the interaction between s_1 and s_2 using unitary evolution of the wavefunction.

$$((\alpha\,|\uparrow\rangle + \beta\,|\downarrow\rangle) \otimes |ready\rangle)_{12} \rightarrow (\alpha\,|\uparrow\rangle \otimes |up\rangle + \beta\,|\downarrow\rangle \otimes |down\rangle)_{12} \tag{3}$$

The final state of s_{12} with respect to s_3 is the superposition $|\psi\rangle_{12/3} = (\alpha\,|\uparrow\rangle \otimes |up\rangle + \beta\,|\downarrow\rangle \otimes |down\rangle)_{12}$. This state is a description of s_3's knowledge of s_{12} rather than a state with ontological meaning attributed to it. Using it, s_3 can predict what they will find if they were to interact with s_{12} themselves. At this point s_{12} has *no definite value* with respect to s_3 even through s_1 has a definite value with respect to s_2 (and s_2 has a definite value with respect to s_1).

This analysis deals with a clear case of measurement. However, RQM claims that values become definite at *any* interaction where systems become correlated, irrespective of decoherence, the creation of records or other signifiers of textbook measurements. Whenever two systems become correlated a value is actualised, even for microscopic systems. Given that all definite values are relative, actualisation does not prevent the measurement of interference effects by other systems since relative to them there is still a superposition.

6.2.2 Ontology

This section will take a closer look at the ontology of RQM. The previous section established that the wave function has no ontological meaning; instead the fundamentals of the theory are the manifestation of definite values of observables at the point of the interaction between systems. In many ways this is similar to the object-property ontology of classical mechanics but with two major differences: a) properties are only attributed at interactions and b) properties are relative. This gives up the classical idea that objects have defined properties at every moment in time (and indeed the idea that objects come before properties). This move is motivated by theorems such as the Kochen-Specker theorem that indicate that not all variables can have definite values simultaneously without violating the predictions of quantum mechanics.

Rovelli (2018) calls the actualisation of a variable at an interaction an *event*. In classical mechanics every moment in time is populated by such an event where a physical system has definite values. In quantum mechanics the ontology is much sparser; Laudisa and Rovelli (2019) call this a *sparse* or *'flash' ontology*. This event ontology can be interpreted in various ways. Early versions of RQM maintain that events are entirely relative to the systems involved and have no ontological status beyond that. Later versions (Adlam & Rovelli 2023; Dorato 2016) take the event ontology as objective. It is an absolute, non-relational fact that an event has occurred and some value has been actualised; but only relative to the systems involved is there any *particular* definite value.[2]

There have been a number of attempts to flesh out this ontology. Candiotto (2017) gives an ontic structural realist interpretation of the theory. Her argument is that the relational nature of RQM is perfectly suited to an ontological interpretation that takes the structure and relations of a theory to be ontologically prior to the objects in it.

On the other hand, Dorato (2016) rejects the possibility of a relation without the existence of its relata. He also raises the question of how to address non-interacting isolated systems (and systems between interactions). To answer these worries he suggests *dispositionalism* where quantum objects exist without intrinsic properties but rather with "intrinsic dispositions to correlate with other systems/observers

---

[2] One difficulty for RQM is identifying exactly when an interaction has occurred and a value actualised. Rovelli (1998) suggests a solution to this. Interactions and the formation of correlations can be modelled with a probability distribution over the times an interaction might be said to occur. Similarly to how the wavefunction in RQM describes the different possible outcomes but only one becomes definite, this probability distribution describes the possible times of interactions relative to an outside observer but the interaction occurs at one of these times relative to the systems involved. In some cases of weak interaction the probability distribution also includes the possibility that no interaction takes place at all. If we take the occurrence of events as a non-relational fact then it must be presumed that the time of the event (and indeed if there was an event at all) is also non-relational. However, no information about this is accessible to any observers except the ones included in the interaction.

O, which manifest themselves as the possession of definite properties q relative to those Os." (Dorato, 2016, p. 239). This fits well with the original presentation of the theory (Rovelli 1996, van Fraassen 2010) where systems are defined by the set of questions that can be asked of them; i.e. the set of properties that could take a definite value when enquired about.

Metaphysical indeterminacy can also be used to examine how different events relate to each other. Proposed by Calosi and Mariani (2020), the suggestion is that metaphysical indeterminacy can help us explain various problems such as what to say about non-interacting systems. They propose this as an alternative to the other understandings of the ontology, but the views are not mutually exclusive.

An outstanding problem that has not received much attention is what exactly a system is in this theory. The fundamental event ontology is of definite values of variables and not of objects and systems. However, defining systems to which these variables belong must be done prior to the actualisation of definite values between them, else the interaction cannot be derived. RQM largely assumes that there are well defined systems between which interactions take place. Dorato's dispositionalism and Calosi and Mariani's indeterminacy attempt to deal with how to describe systems *between* interactions but this does not address the question of how to define a system in the first place. Adlam and Rovelli (2023) state that events are what is fundamental, and systems are identified by tracking lawlike patterns among the events that pick out systems evolving over time. However, this idea has not so far received much attention or critical development. Adlam (2023) further explores what a system is in relation to how RQM could be reduced to quantum field theory (QFT). In relativistic approaches, systems and trajectories are approximate and emergent only; Adlam briefly discusses some potential solutions to this such as accepting a vague ontology, replacing systems with regions of spacetime, or adding a sharply defined ontology by hand. All of these have their drawbacks (see Adlam for details) and none have been developed in depth.

In addition to the issue of reduction to QFT, we can see further problems for defining systems when it comes to subsystems and complex composite systems even within non-relativistic quantum mechanics. It has not been made clear whether events are only microscopic: either there are events corresponding to macroscopic systems as well, or macroscopic events are built out of sets of microscopic events. In standard quantum mechanics macroscopic systems are commonly described using a tensor product of their subsystems. However, at times Rovelli and other authors deny that we can straightforwardly make sense of tensor product states in RQM and rely on this to solve various challenges (for example see Cavalcanti, Di Biagio, & Rovelli (2023) where this is used in rejecting a potential no-go theorem for RQM). A product state is just a collection of all the information that a third party holds separately about a set of subsystems, but it ignores all the values that the subsystems must have relative to each other. We can only understand systems in isolation by looking at these internal values, tensor products cannot do this. What this means for thinking about macroscopic interactions and the corresponding events has not been made clear. At many other points in the literature there are examples used that clearly involve macroscopic, composite systems and these are described as single events without considering what they are made up of. This potentially leads to a double set of events, one at the macroscopic level and one at the microscopic, for the same interaction. No composition relation for events has been spelled out, and without doing so how we should think about systems remains unclear. Therefore this issue remains a significant outstanding problem for spelling out the ontology of RQM. It is especially relevant for how to think about decoherence in RQM as traditional analysis of decoherence involves thinking about open systems and how correlations form between many subsystems.[3] As decoherence will play an important role in the arguments to come, my conclusions will make some suggestions towards what we should look for in a solution but leave this issue ultimately unsolved.

6.3 Temporal Becoming in RQM

With this overview of RQM in mind and taking an event ontology as the basis for understanding the theory I will now look at how temporal becoming as a model of time's flow fits with this interpretation. The event ontology of RQM makes

---

[3] Adlam (2023; 2024) and Adlam and Rovelli (2023) give some discussion of decoherence and how it runs into issues such as cross-perspective links (see section 6.4), but no clear solution is reached.

it a good fit for temporal becoming and this project has already been begun by Dorato (2016) with a comparison to his version of localised temporal becoming. RQM's formulation in terms of information offers further resources to explore the link between records and the open future that the previous chapters have focused on. This would be an extension of Dorato's view of becoming in RQM and shows that RQM supports the use of records in conjunction with temporal becoming.

First I will summarise the overall picture of temporal becoming developed in previous chapters (and how it differs from Dorato's account). Then I will work out the details of how this picture of temporal becoming would work in RQM.

6.3.1 Temporal Becoming and Open Future Indeterminacy

Becoming in time is defined as events *becoming definite* or *coming into existence*; also commonly described as the *occurrence* or *happening* of events. The particular version of the view that I am using is a more recent attempt to use becoming to reconcile time's flow with a block universe. However, I argued that we should instead focus on the perspectival aspect of this account where the tension between presentism and eternalism comes from a failure to specify the position of the observer in relation to the time they are describing. When situated at a moment inside of time, the events at that moment are all that exist and past and future ones do not - this is presentism. But when you remove the perspective of any specific observer you can take a god's eye view of the entirety of time with all the moments taken together -this is the eternalist view. An event coming into existence at a specific moment is temporal becoming, and the change in existence from moment to moment in the presentist perspective constitutes the flow of time.

Dorato (2016) - in his attempt to connect localised temporal becoming to RQM - claims temporal becoming has three necessary ingredients: 1) events (an event ontology), 2) local succession of events along a worldline, and 3) irreversible succession.

Condition (3) comes from a prevalent assumption that once something has Become (become determinate) it remains so. This has been an assumption right from the beginning of discussions of temporal becoming which used a growing block ontology. When applied to localised temporal becoming, the same principle has been assumed to hold: once an event has become it remains definite thereafter in the future. I have argued in previous chapters that this assumption has no basis. An event comes into existence and at the next moment goes out of existence as the next event comes into existence. This occurs both forwards and backwards in time. The irreversibility of becoming comes from the desire to be able to talk about a definite past that, once is *has become*, is fixed even if it no longer exists relative to the present.

Related to this is the question of the indeterminateness of the future. Dorato does not attempt to preserve an open future and thinks that, although they have not yet become relative to the present, future events are fixed by the eternalist perspective. However, I have argued in previous chapters that this ignores the presentist side of this combined view in favour of a more familiar eternalist view. Doing so weakens the concept of becoming and leaves out essential features of the structure of reality. As such I claim that the future is metaphysically indeterminate with respect to the present. This is because the localised present state does not determine what the future will be. Outside of quantum mechanics this is due firstly to the lightcone structure that limits what can be known (and therefore used in predictions) in a localised present, and secondly due to higher level probabilities in theories such as thermodynamics (which are potentially but not necessarily derivative from quantum probabilities). Both of these are examples of ways in which the amount of information available within the present perspective is limited and prevents the future from being predicted with certainty.

If this notion of indeterminacy describes the openness of the future then we are left in need of an account of the determinateness of the past; this is where epistemic considerations can be most fruitful. To restore a determinate past I have argued that determinacy can be based on records of the past that are preserved in the present. Records are physical systems that provide information about the past. This information preserves the determinacy of the past beyond the moment it came into being; no such information is available for future events so an asymmetry is created. The physical creation of records plays an active role in physics, illustrating that the epistemic asymmetry is an integral part of the way we describe the world as well as how we identify temporal asymmetry within physics.

To summarise, Dorato's three ingredients for localised temporal becoming can be amended to: 1) events, 2) local succession of events on a worldline, and 3) asymmetric succession based on records of the past and an open future. Condition 3 does a huge amount of work in establishing temporal becoming as an ineliminable feature of reality.

6.3.2 RQM: Temporal Becoming

With this picture in mind I now turn to indeterminacy and temporal becoming in RQM. This section and the next will respectively explore the general picture and then future indeterminacy. Establishing a definite past on the basis of records will then be the topic of the rest of the chapter.

The first of the amended conditions for temporal becoming - the existence of events - can be instantly fulfilled with RQM's event ontology. The actualisation of a value at an interaction corresponds to the becoming of that event. The second condition is also fulfilled. Becoming occurs along the worldline of each system and tracks the succession of interactions that it is involved in. Dorato describes becoming in RQM as follows:

> "In the form of relativistic becoming endorsed by RQM what we have is a crisscrossing of little ripples, unrelated to each other, which give us local, non-worldwide becoming (corresponding to the incomplete information that each observer has about the universe, given that she is inside it). The fact that in RQM we have no universal and cosmic tide of becoming also corresponds to the locality of RQM" (Dorato, 2016, p. 260-261)

(Note that the locality of RQM has been challenged and discussed at length. Martin-Dussaud, Rovelli, and Zalamea (2019) discuss the senses in which RQM is and is not local. The relevant sense in this context is that interactions are localised between systems, and events can be tracked along the worldline of a single system. Non-locality in EPR experiments is debated - though the above reference argues that there is no *meaningful* sense of non-locality in RQM - but will not be considered here.)

The ontology of RQM also fits well with the perspectival relation between presentism and eternalism. This is perhaps unsurprising as Rovelli is one of the proponents of this view (Rovelli 2019), although he has never explicitly connected this view of time to RQM. The presentist perspective is made up of the definite values at events and the epistemically understood wavefunctions that each system uses to encode information it has gathered from past interactions and to predict future definite values. The set of all events is objective and forms the eternalist perspective. Note that the eternalist perspective does not include the particular definite values relative to the systems involved at each event, only that there is an event. So this is a clear example of how the eternalist block does not provide a complete description of the world and must be supplemented by the presentist perspective.

6.3.3 RQM: The Open Future

The third condition regarding an indeterminate future and a definite past provided by records will take more exploring.

Metaphysical Indeterminacy in RQM is discussed by Calosi and Mariani (2020). They base their account on Wilson's (2013) determinable account of metaphysical indeterminacy; the variables of systems are determinable properties and the values that are actualised in events are the determinates.

Calosi and Mariani (2020, pg 162) adapt the general definition for quantum metaphysical determinacy to apply it to RQM as follows:

> "**Metaphysical Determinacy**. A quantum system s_1 is *metaphysically determinate* with respect to observable O iff, necessarily, for every other quantum system s_2, s_1 has a unique value v_2 of O relative to s_2. Or, equivalently: necessarily, s_1 has O = v_2 relative to every other quantum system s_2."

There is a general point to be made about the way that Calosi and Mariani consider metaphysical indeterminacy. Their definition considers the absolute status of a variable of a system as determinate or indeterminate; only if *every* other system has a definite value of a variable will that system be metaphysically determinate with respect to that variable. As such, a system cannot have a determinate variable unless it has interacted with every other system in the universe (and interacted in way that actualised that specific variable at that time). This is impossible; every system comes out as metaphysically indeterminate with respect to all its variables (Calosi and Mariani note this in footnote 35 of their paper). This is not a huge problem as only relational values are important in RQM. However, given that relationism is the foundation of RQM, it seems straightforward that we should also consider indeterminacy relative to perspectives. Calosi and Mariani do not give much attention to this. They only take into account indeterminacy due to the combination of different values being taken relative to different systems. In terms of the perspectival relationship between presentism and eternalism defined previously, this corresponds to indeterminism in the eternalist perspective.[4] This neglects what metaphysical indeterminacy can be found within a single perspective (relative to a single system) that corresponds to the presentist view. Values become definite (i.e. determinate in some way) at the interaction between systems; a single value is actualised and becomes definite relative to the other system involved in the interaction (but not to outside observers). Therefore, it is important to consider determinateness from the perspective of a single system so as to satisfy RQM's fundamentally relational nature.

This is easy to do, we already have all the necessary tools. It simply requires understanding determinateness as a perspectival property rather than an absolute one. The definition of metaphysical determinateness can be adapted by simply removing the references to multiple systems:

> **Metaphysical Determinacy:** quantum system s_1 is *metaphysically determinate* with respect to observable O relative to system s_2 at a time t if system s_1 has a unique value v_2 of O relative to s_2 at time t.

A system being metaphysically determinate relative to another system is nothing more than for a definite value of a variable to be actualised in an interaction

---

[4] This already differs from the standard eternalist view such as that of Dorato in that the eternalist view of events contains no determinate facts at all. In RQM all determinate values are relational.

between the systems. During an interaction a correlation is formed between two systems that acts as the physical determinate for the determinable property. Prior to interaction no value can be assigned to s_1 by s_2 and so s_1 is indeterminate with respect to s_2. The actualisation of the value when the systems become correlated corresponds to temporal becoming. Becoming is localised and is part of a succession of interactions relative to a specific system. All events that are not part of this succession are events that have not (yet) occurred (or become). In the classical literature this refers to just future events. But in the localised version of becoming it includes not only the future but all events not on the worldline of the observer, in other words, all systems that the observer has not interacted with.

There is an important clarification to note here. Most often in discussions of quantum indeterminacy the indeterminateness is seen to come from superpositions. For example, Wilson (2013) argues that a spin superposition of up and down values is a case of metaphysical indeterminacy, as the determinable of spin has no single determinate (only a superposition of them).[5] This is not the case in RQM. No superpositions are ever part of the ontology; only single values are actualised. Superpositions, as part of the wavefunction, are only a predictive tool that represent what information has been gained in the past and hence cannot be seen as a case of metaphysical indeterminacy in and of themselves. The superpositions make predictions about the future but open future indeterminacy is due to the lack of knowledge prior to interaction. In this bigger picture, this also means that quantum indeterminacy in RQM is no different from indeterminacy in areas such as thermodynamics that I laid out in chapter 3. Here the indeterminacy was perspectival due to limits from both the lightcone and what information is available in the macrostate. While this is potentially grounded on quantum probabilities, higher level probabilities are often seen as less fundamental and of a different type. In RQM however, quantum probabilities and indeterminacy are equally due to limits on knowledge due to your perspective and there is no substantive difference.

---

[5] This is already contentious in the context of - for example - an Everettian view of quantum mechanics or any interpretation that adopts some sort of state realism. In this case the wavefunction is seen as a ray in Hilbert space and there is no indeterminacy. It should not be read as being indeterminate between two measurable properties at all but as a single determinate quantum state.

So far RQM is very hospitable to localised temporal becoming. It can provide events, local succession and an indeterminate future. What is more, the indeterminacy of the future - and the corresponding becoming along a worldline - is an integral part of the theory. Section 6.2.1. laid out a simple situation where s_1 has a determinate value relative to s_2, but relative to s_3 both are in a superposition with no determinate result. The indeterminacy with respect to s_3 is necessary to preserve the interference effects that we expect to find in quantum mechanics and to make correct predictions about the future. If s_3 were to make a measurement of interference effects then its description of s_{12} in a superposition state is an accurate predictor of the measurement results s_3 would find. The indeterminateness from s_3's perspective is as important as the determinateness from s_2's perspective; without it the empirical results of quantum mechanics (finding interference effects) could not be explained. Hence metaphysical indeterminacy relative to a perspective is essential for a complete explanation of reality.

The remaining aspect to consider is the definite past. RQM is fundamentally time symmetric (Adlam & Rovelli 2022; Di Biagio, Dona, & Rovelli 2021) in its mathematics and its ontology (from an eternalist perspective). It appears asymmetric when we consider the wavefunction that contains information about past interactions and is used to predict future ones.[6] At first glance this seems to highly motivate the picture that I have suggested, in which the definite past comes from having information about the past in the form of records. The wavefunction represents a 'record' of the past interactions where values have become definite; this is physically realised in the correlations between systems. However, the theory cannot be straightforwardly interpreted in this way as it does not use a robust sense of records. Rovelli stresses that values become actualised in *any* interaction regardless of whether or not it is a measurement and produces physical records. Any correlation counts.

The next section will explore in more detail how the past is made definite and show that some current problems for RQM relate to this issue of how definite values from

---

[6] Rovelli claims that the wavefunction is essentially used to map between events in a lawlike way and that this can be done in either direction in time. Hence the wavefunction is asymmetric but only trivially so. The question of how to talk about the initial state of the universe is not addressed in the RQM literature but would likely determine which way to apply the wavefunction.

the past are preserved. The conclusion from this will be that while it is coherent to maintain that values become definite at any interaction, the continuity of this definiteness into future interactions relies on the presence of robust physical records.

6.4 The Propagation of Definite Values

Having laid out the general connection between temporal becoming and RQM, I now turn to the status of past actualisations and examine what happens to them as time passes. This touches on some recent issues that have been discussed in the literature on RQM relating to how systems share information over time. I will show that applying the idea that past values remain definite only if physically robust records of them persist into the present can help to understand these problems and reduce their impact. Some aspects of these issues remain and this chapter will not attempt to resolve them entirely. However, it will show that, to the extent that RQM has been developed so far, the most promising avenue for understanding it is to assume a definite past only in the decoherent domain based on the presence of records. The first problem I will consider is what happens in interactions where no basis is selected and multiple possible variables seem to be actualised. Then I will discuss a Wigner's friend type situation that has been raised as an issue for RQM.

6.4.1 Interactions Without Outcomes

Rovelli claims that values are actualised in any interaction:

> "Value actualisation happens at interactions because variables represent the ways systems affect one another. Any interaction counts, irrespective of size, number of degrees of freedom, presence of records, consciousness, degree of classicality of S', decoherence, or else because none of these pertain to elementary physics" (Rovelli, 2018, p. 5)(where S' is the system interacting with the primary system S).

However, this calls into question what exactly is being actualised and what a quantum event really is, particularly in cases of microscopic interactions that can take place without any basis selection occurring.[7] The examples commonly used to ex-

---

[7] Similar points are raised in Muciño, Okon & Sudarsky (2021)

plain RQM use clearly defined bases and variables irrespective of the actual details of the systems involved. This is not reflective of the complexity of quantum mechanics. To see why we will consider Zurek's (1982) presentation of measurement and decoherence that I will lay out in some detail. The full complexity of the situation illustrates exactly what ways RQM struggles to capture the full scope of quantum mechanics.

Zurek's account of decoherence explicitly links measurement to the creation of records through interaction with the environment - although this applies to any interaction where decoherence might arise and not just formal laboratory measurements.[8] He shows that there are two stages to a measurement: first, the interaction between the system and a measuring system and second, an interaction between these two systems and the wider environment.[9] To make this explicit he considers the case of a pair of two-state systems interacting. The spin in the z direction of one system can be described by the $|\uparrow\rangle$, $|\downarrow\rangle$ basis. The other system is doing the measuring and can be described in the basis of 'ground' $|\eqcirc\rangle$ and 'excited' $|\circeq\rangle$ states. Both systems start in a maximally entangled state and they can be written in many different orthonormal bases. For example spin can be rewritten as $|\odot\rangle = (|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$, $|\oplus\rangle = (|\uparrow\rangle - |\downarrow\rangle)\sqrt{2}$ or as $|\rightarrow\rangle = (|\uparrow\rangle + i\,|\downarrow\rangle)/\sqrt{2}$, $|\leftarrow\rangle = (|\uparrow\rangle - i\,|\downarrow\rangle)\sqrt{2}$. Similarly the atom basis can be rewritten $|+\rangle = (|\circeq\rangle + |\eqcirc\rangle)/\sqrt{2}$, $|-\rangle = (|\circeq\rangle - |\eqcirc\rangle)\sqrt{2}$ or as $|\perp\rangle = (|\circeq\rangle + i\,|\eqcirc\rangle)/\sqrt{2}$, $|\top\rangle = (|\circeq\rangle - i\,|\eqcirc\rangle)\sqrt{2}$.

The interaction between these two systems transforms the initial state

$$|\psi_i\rangle = ((|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}) \otimes ((|\circeq\rangle + |\eqcirc\rangle)\sqrt{2}) \tag{4}$$

to the final state

$$|\psi_f\rangle = (|\uparrow\rangle \otimes |\circeq\rangle + |\downarrow\rangle \otimes |\eqcirc\rangle)/\sqrt{2}) \tag{5}$$

---

[8] This repeats in more detail the analysis laid out in the previous chapter regarding whether a single photon scattered off a central atom can be regarded as a record. However, this situation does not require us to accept quantum darwinism as a theory of decoherence. The example is deliberately chosen from early in the development of decoherence to illustrate that these are basic and widely accepted results independent of any specific program.

[9] The 'measuring system' Zurek uses is a very thin notion consisting of just a two state system - it does not imply a measuring device that a complex human observer would use. For complex measuring devices the device on its own can often play the role of both apparatus and environment.

where the interaction Hamiltonian is

$$H^{AS} = g(|\perp\rangle \langle\perp| - |\top\rangle \langle\top|) \otimes (|\uparrow\rangle \langle\uparrow| - |\downarrow\rangle \langle\downarrow|) \tag{6}$$

and acts over the interval $\tau_i = \pi\hbar/4g$ (where g is the coupling constant).

In RQM this interaction would give events with either $|\uparrow\rangle_{spin/atom}$ and $|\circeq\rangle_{atom/spin}$ or $|\downarrow_{spin/atom}\rangle$ and $|\eqcirc\rangle_{atom/spin}$ as possible definite values with equal likelihood (where the values are relative to the systems involved, as indicated by the '/' in the subscript). This event would occur probabilistically sometime over the period τ_i.

However, Zurek shows that this final state can just as easily be rewritten as

$$|\psi_f\rangle = (|\odot\rangle \otimes |+\rangle + |\oplus\rangle \otimes |-\rangle)/\sqrt{2}) \tag{7}$$

which in RQM would give an event with either $|\odot\rangle_{spin/atom}$ and $|+\rangle_{atom/spin}$ or $|\oplus\rangle_{spin/atom}$ and $|-\rangle_{atom/spin}$ as possible definite values with equal likelihood.

Zurek goes on to show that only with the further interaction with the environment does a basis get singled out and a record formed (the singling out of a basis fixes what states of the atom correspond to what spin states, hence allowing the atom to act as a record). This is the process of decoherence where the off-diagonal terms in the density matrix of the system (when the environment is traced out) tend towards 0.

Under Rovelli's account the first stage of Zurek's measurement is an interaction and hence results in the actualisation of a definite value.[10] But what value is it? The interaction has resulted in an exchange of information between the systems such that the variables of one are correlated to the variables of the other. But no single variable is singled out and we can freely switch between different bases. Each possible basis defines a different possible set of events with different values actualised. Adlam and Rovelli suggest the following to describe what happens here:

---

[10] It is essential in RQM that this sort of interaction *does* produce an event so as to produce a well populated flash ontology that describes all levels. Without events from these microscopic interactions it would seem that microscopic systems couldn't be said to exist at all. This is presumably an unwelcome consequence.

> "quantum events do not typically have the simple form 'variable V taking a value v relative to Alice.' Rather they must have a conjunctive form: 'variable V_1 taking value v_1 relative to Alice, and variable V_2 taking value v_2 relative to Alice,. . . ' and so on, specifying definite values for each of the variables singled out by the interaction Hamiltonian in all of the different possible bases for it." (Adlam & Rovelli, 2023, p. 15)

This must then be extended to account for the different bases of the measuring system that take a value relative to the system being measured as well. For the situation above, an event for the atom would be the actualisation of (for example) $|\uparrow\rangle_{spin/atom}$ and $|\odot\rangle_{spin/atom}$ and so on for all possible bases. For the spin system the event would be (for example) $|\circeq\rangle_{atom/spin}$ and $|+\rangle_{atom/spin}$ and so on.

This seems to violate the common textbook interpretation of quantum mechanics that generally states that systems cannot have definite values of all variables at once. Having a definite value of one variable should mean that the value of a second, non-commutable, variable is indeterminate with a probabilistic distribution over the possible values that can be calculated. This, however, is only a problem when considering sequences of measurement. Measuring one variable - for example spin in the x direction - gives a definite value of that variable. A subsequent measurement of spin in the y direction should give up and down with equal probability. At a single point multiple variables can take on definite values; what needs to be ensured is that in subsequent interactions these values are not preserved and only probabilistic predictions apply. This will mean that one cannot, in successive measurements, get results for spin in the x, y, and z directions together.

The question then, is what we use to predict the next measurement in the sequence. The wavefunction in RQM is a predictive tool that encodes information about past interactions and predicts the outcomes of future ones. When an event occurs and a value becomes definite this updates the wavefunction, reducing it to the eigenstate for that value (which looks like 'collapse'). Applying this to the situation above of Zurek's first stage of measurement, what should the wavefunction be updated to? A collapse in one basis updates the wavefunction to $|\uparrow\rangle_{spin/atom}$ but in another basis the collapse gives the wavefunction $|\odot\rangle_{spin/atom}$. These are orthonormal bases related by $|\odot\rangle = (|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$. As such, these are incompatible wavefunctions; the system cannot be described by both. Taking this further, we can ask which wavefunction should be used to make predictions about future interactions (for example a second interaction between the spin-atom systems on a short enough time span that the wavefunctions have not appreciably spread). The two wavefunctions give incompatible predictions (either $|\uparrow\rangle_{spin/atom}$ with certainty or $|\uparrow\rangle_{spin/atom}$ and $|\downarrow\rangle_{spin/atom}$ with equal probability). So as not to violate the uncertainty principle between non-commuting variables, consistency is required between what the interaction updates the wavefunction to and what is used to predict subsequent measurements. Hence if the interaction gives $|\psi\rangle_{spin/atom} = |\odot\rangle_{spin}$ then $|\odot\rangle = (|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$ must be used to predict future interactions in the $|\uparrow\rangle$ / $|\downarrow\rangle$ basis. But given that it is not clear what we should update the wavefunction to, there is nothing to ensure that this consistency holds.

What this makes clear is that the ontological solution - expanding the notion of event to include the actualisation of multiple variables[11] - does not solve the problem entirely and we are left with questions about the wavefunction and how to make predictions. Adlam and Rovelli summarily dismiss such concerns by saying that standard quantum mechanics does not have anything to say about how to make predictions when qubits interact, because this does not count as a measurement as we would understand it. Quantum mechanics only makes *predictions* about outcomes at a decoherent level. But this undermines the very basis of RQM where the wavefunction is understood to be a summary of information from past interactions used to predict future ones. If there is no meaning to prediction at the level of qubits then this dramatically undermines RQM. At the very least RQM must provide an answer to an additional aspect of the situation that Zurek demonstrates: if the interaction continues for an additional time $\tau_r = 3\pi\hbar/4g$ then the final state (Eq.4 or Eq.6) is transformed back into the initial state (Eq.3). No information about any specific outcome at τ_i is preserved at τ_r. A second interaction between the two systems would proceed with the original pre-interaction probabilities. In

---

[11] It is somewhat unclear whether this should be taken as a single conjunctive event where multiple variables become definite, or whether the original definition of an event as a variable becoming definite should be preserved and the situation interpreted as many overlapping events (in the same place and time). However, this difference should have little effect on our understanding of RQM.

RQM the wavefunction would be updated at τ_i and evolve from there, presumably spreading out again under evolution according to the Schrödinger Equation. RQM has not answered the question as to whether at τ_r the predictions concerning a second interaction of the two systems made by the updated wavefunction match those from unitary quantum mechanics.

In light of this we can use the model of temporal becoming previously developed to suggest a solution. This model says that definiteness of the past is dependent on records being propagated into the future. Temporal becoming describes how things are made definite at each moment (or event) but this does not guarantee that this definiteness is propagated into the future. For this to be the case something physical - a record - must be passed forward in time. Only when a robust record is created should we update the wavefunction that is used for predictions.

As already noted, this fits well with RQM's principle that information is contained in physical variables. This is also what the problem described in this section comes down to: which variables are correlated in informative ways and what this information actually is. The interaction between the two systems establishes a correlation between them. In Rovelli's information-theoretic sense this means that information has been passed between them and this is encoded in the physical correlation. This is the basic idea on which RQM was developed. However, in actuality the correlation is not in any particular basis and while the states of the two systems depend on each other, it is not a strong enough dependence for the state of one system to tell us anything definite about the state of the other. This assumption of strong dependence seems built into RQM, which requires that the correlation makes definite a value of a specific variable in a specific basis such that the state for one system is tied to a particular state for the other system. This goes beyond the information that the physical correlation actually contains.

Ontologically there may be good reason to argue that any information exchange should be counted as an event - a correlation has been created and there is always a thin notion of information related to this. However, a further implicit assumption is being made in RQM, which is that all events (and the specific values made actual in them) should be counted when predicting the future and all events should be used to update the wavefunction. This additional step is unwarranted in situations - such as the one presented above - where no physical records are made. Only a minimal amount of information is physically preserved forwards in time and only this should be used to affect future predictions. The best description for making predictions is still the original superposition that has not been updated based on values from the event.

Before moving on I will show how this fits into the picture of metaphysical indeterminacy that I have developed. I have argued that future interactions and systems that have not interacted can be treated as metaphysically indeterminate (from the perspective of a single system) in the same way that we can treat the open future as indeterminate. We can expand this to include past systems for which no physical records are made and propagated into the future. While there *was* a definite value (or set of values) actualised at the time of the interaction, at a future time this value no longer exists and the value is indeterminate. This is in the same way that future interactions *will* actualise a definite value but *currently* the variables of the system are indeterminate. Maintaining multiple different values into the future leads to incompatible predictions. As such, it must be the case that none of these values are propagated forwards or used in predictions. Only in further stages of interaction with the wider environment – what Zurek calls the second stage of measurement – are records made and a pointer basis is selected. At this point, information about a specific basis is encoded into the physical correlation between systems, and the definite value is maintained into the future. The system can continue to make definite claims about what value was actualised in the past interaction and use this to make predictions of future interactions.

6.4.2 Intersubjectivity and Cross-Perspective Links

There is another place where comparing results at different times causes difficulties for RQM. This concerns what happens when multiple different systems make a sequence of measurements and try to communicate their results to each other (similar to the Wigner's friends thought experiment but applying to both microscopic and macroscopic systems). This issue has recently been discussed in the literature (Adlam & Rovelli 2023; Brown 2009; Pienaar 2021). I will present the problem and the proposed solution from Adlam and Rovelli as well as some limitations of it. I will show that, despite outstanding problems with the ontology that prevent a full explanation of decoherence in RQM, the situation so far strongly motivates the same conclusion: that definite values should not propagate over time in sequences of interactions (events) at the microscopic, non-decoherent level but should when decoherence takes place and records are made.

To illustrate the problem I will return to the initial example presented in section 6.2.1 and extend the situation further. The basic situation was three systems. System s_2 performs a measurement on system s_1 and gets a definite value as an outcome. System s_3 has only the starting states of s_1 and s_2 and therefore describes the measurement interaction as a unitary evolution resulting in a superposition of outcomes. In the original presentation of the theory (Rovelli 1996) this leads to an interesting consequence if s_3 physically interacts with s_{12} after the first measurement of s_1 by s_2. This measurement would proceed in exactly the same way that the interaction between s_1 and s_2 did and s_3 would find $|\psi\rangle_{12/3} = (|\uparrow\rangle \otimes |up\rangle)_{12}$ or $|\psi\rangle_{12/3} = (|\downarrow\rangle \otimes |down\rangle)_{12}$ with probabilities α^2 and β^2 respectively. This leads to a situation where we could have $|\psi\rangle_{1/2} = |\uparrow\rangle_1$ as the result of the first measurement of s_1 by s_2 and $|\psi\rangle_{12/3} = (|\downarrow\rangle \otimes |down\rangle)_{12}$ as the result of the second measurement of s_{12} by s_3. In RQM, these contradictory results between different interactions is a possible physical situation resulting from the information held by s_3. Before s_3 measures s_1, s_1 has no physical values with respect to s_3 regardless of what physical values it may have with respect to other systems.[12] This, however, is purely from the perspective of s_3. When considered from the perspective of s_2 there is a worrying result. s_2 measures s_1 and initially finds $|\psi\rangle_{1/2} = |\uparrow\rangle_1$; but the subsequent interaction with s_3 overwrites this value. Any communication between s_2 and s_3 is an interaction that overwrites s_2's original observation and leaves s_2 unable to communicate their original value to s_3. s_3 will find s_2 reporting up or down with equal probabilities (according to the information they hold of s_{12}) regardless of the original result.

---

[12] Note, however, that there are still the normal consistency conditions such that we will never get a situation where there is inconsistency within in a single interaction. System s_3 will never find $|\psi\rangle_{12/3} = (|\uparrow\rangle \otimes |down\rangle)_{12}$.

That a solution is found for this problem of sequences of interactions is important for two (interconnected) reasons. First, to avoid absolute solipsism. There would be a problem with the theory if a subject is unable to communicate their own experiences. Second, leading on from this, a lack of intersubjective communication is not just a worry for what it would mean for our intuitive understanding of how we can communicate, but it would also pose a problem for the methods of empirical confirmation that are essential for science. If we were truly unable to communicate our own knowledge and to use our own subjective experiences as representative of objective reality then the methodology of scientific enquiry and the development of quantum mechanics itself as a theory would be impossible. To empirically discover a theory that rules out the possibility of empirical confirmation is a contradiction (See Adlam (2022) for more on this problem).

Adlam and Rovelli have recently considered this problem and suggested a new postulate for RQM that eliminates it. They propose a *cross-perspective link*:

> "**Cross-perspective links:** In a scenario where some observer Alice measures a variable V of a system S, then provided that Alice does not undergo any interactions which destroy the information about V stored in Alice's physical variables, if Bob subsequently measures the physical variable representing Alice's information about the variable V, then Bob's measurement result will match Alice's measurement result." (2023, p. 5)(where in this scenario Alice would be s_2 and Bob s_3)

This is further explained as follows:

> "a measurement on Alice aiming to establish her information about the variable V should be understood not as probing her instantaneous state at the time but as 'looking back' in a nonlocal way at the value that becomes definite in the most recent interaction." (2023, p. 8)

This solves the problem above by straightforwardly stipulating that the results of the second measurement of s_{12} by s_3 will automatically match the result of the first measurement of s_1 by s_2. An interaction is seen as looking at the recent history of *both* systems involved to determine the outcome. The probabilities of the results of s_3's measurement remain the same as the wavefunction used to describe s_{12} is still in a superposition and this determines a real probability with respect to s_3. The value s_1 is only definite with respect to s_2. But given that s_{12} is as equally involved in the measurement as s_3 is, its history and the value actualised in the measurement by s_2 also affect the outcome. Furthermore, Adlam and Rovelli argue that s_3 would still be able to detect interference effects if the right experiment was performed. In RQM all knowledge is physical and hence the superposition in $s_{12/3}$ is physical and the interference in it can be measured. Interference measurements by s_3 can be seen as destroying the information about the previously actualised value in the measurement by s_2 such that s_2 can no longer use it in predictions.

There are a number of doubts one might have about this solution, including whether it can genuinely deal with interference effects as they claim; I will not lay out all possible objections here. But the solution is immediately troubling. The non-locality of 'looking back' is not explained; introducing a new form of non-locality into quantum mechanics is not appealing. Additionally 'looking back' is only done during measurements in the same basis as the previous measurement; when s_3 measures interference in a different basis the previous results from s_2 are disregarded. This difference seems arbitrarily contrived purely to get the results needed.[13] But putting aside these concerns I will note that if the solution from the previous section is adopted - that interactions without a preferred basis are not used for future predictions - then the immediate problem disappears in non-decoherent cases as s_2 does not get a unique definite value to use in predictions and instead uses the superposition state to make predictions (just like s_3). No continuity is expected between the original and subsequent interactions so the overwriting of s_2's result is not worrying, and the detection of interference effects in any basis is straightforward. The cross-perspectives problem still remains in cases where decoherence *has* occurred and s_2 *does* get a single definite value (and expects a continuous experience), but this is a much restricted set of the original problem (and there are no problems with interference effects).

Furthermore, the concerns that motivated cross-perspective links are only relevant

---

[13] This is confirmed further in Adlam (2023) where she notes in a footnote the ad hoc nature of the postulate with the hope that further work will show that it can be derived naturally from decoherence, although there are doubts about whether this can be achieved as discussed below.

to the macroscopic regime; despite the postulate applying universally. Adlam and Rovelli note in several places that we should not try to understand what it means for a qubit to experience a definite value; they clearly do not have subjective experiences comparable to our own human consciousness. As such, neither solipsism nor a failure of intersubjective communication should be worrying. These features of experience, along with the expectation of continuity over time where values are preserved in future interactions, belong to a much richer notion of observer than RQM uses.

To look for a solution to the cross-perspective links problem that takes into account the differences between observers at different levels, we should examine how observers are connected to correlations - which is the foundation of RQM. Any system is an observer, regardless of scale or complexity, even a single qubit. This follows from the principle that information is physical and any physical correlation in the world represents an exchange of information; associating an outcome with any correlation is the major interpretive step made to deliver this principle. However, not all correlations are the same. A robust macroscopic correlation is very different from the correlation between two microscopic systems. This suggests that it may be possible, and indeed strongly in line with RQM's principles, to interpret different correlations in different ways.

Different correlations represent different standards of information and the features associated with a more complex observer should be associated with the sort of correlations that these sorts of observers enter into, and not assumed for more basic correlations. The work on decoherence in the information framework, laid out in the previous chapter, makes this especially clear. Decoherence describes the emergence of correlations that are robust and impervious to many interactions with different observers. These correlations are said to emerge classically precisely because they exhibit features that purely quantum correlations do not. They have been labelled as records because they latch onto what we mean when we say that information is preserved. Decoherence, at first glance, has all the indications of presenting a solution to the cross-perspectives problem and explaining how definite values are propagated over time. All of the concerns raised about subjective experiences, intersubjectivity and scientific measurement apply within the decoherent domain.

And versions of decoherence such as quantum darwinism explicitly argue that one of the central effects of environment-induced decoherence is to produce intersubjectivity by creating records that allow for objective agreement among multiple observers about the state of the system.

However, there are a number of difficulties with deriving decoherence from the current version of RQM and, ultimately, work on decoherence cannot progress until the ontology is worked out in more detail. The first problem is with the use of tensor product states. Thinking back to Zurek's (1982) account of decoherence laid out in the previous section, when he comes to modelling the second stage - interaction with environment to produce decoherence - he does so using the combined state of the system-environment written as a tensor product. But RQM, as previously mentioned in section 6.2.2, can only understand tensor product states as states relative to a third-party system (because tensor products cannot account for what states the different components have relative to each other). This means that the systems involved cannot use tensor products to derive decoherence *relative to themselves* the way that standard quantum mechanics does. Adlam and Rovelli (as well as Adlam 2024) argue that we have to assume cross-perspective links at a microscopic level to work out how all the component subsystems coordinate and come to agree on what value is actualised. But if we recognise that we should not expect continuity at the microscopic level this becomes significantly easier - the values relative to individual sub-components can be overridden by future interactions with no difficulties. Decoherence should not be understood as all the subsystems getting together and reaching a consensus on what value is actualised; this implies a level of self-awareness and organisation that should not be attributed to qubits. Decoherence is a phenomenon of emergence where the microscopic level is (conditionally) irrelevant to the behaviour of the system as a whole. This removes the need for cross-perspective links at the microlevel and minimises the concerns about using tensor products.

The real problem with working out how decoherence progresses in RQM is how to think about the ontology of composite systems. As discussed in section 6.2.2, it has not been specified whether interactions between macroscopic systems correspond to distinct events in the ontology or whether these have to be constructed out of microscopic events. The former solution would make modelling decoherence considerably easier; we can see macroscopic systems as patterns among macroscopic events in the same way that microscopic systems are patterns among microscopic events. Composition relations could be identified by looking at correspondence between the different sets of events. This is certainly not a worked out solution but provides a starting point. Adlam (2024) alludes to a solution like this but does not connect the problem to ontology - she merely discusses the possibility of macroscopic results not being derivable from microscopic ones.[14] If we reject macroscopic events as part of the ontology the problem is significantly harder and it is not clear how to proceed.

To fully assess these options and how they would work in RQM would be a far larger project than is possible here and would need to take into account the full range of situations involving complex systems that RQM has to deal with. Without this further work specifying the ontology, a full solution to the cross-perspective problem cannot be developed. However, focusing on how correlations differ at different levels, and being careful about what expectations we have about observers associated with different kinds of correlations, is a useful guide to make progress on this that remains in line with the principles that RQM is based on.

6.5 Conclusion

This chapter has examined how temporal becoming could work within the framework of RQM. In so far as the interpretation has been developed up till now, there is good reason to think it supports localised temporal becoming as a model for time's flow. The event ontology of RQM is a good foundation for quantum temporal becoming, as argued by Dorato, as it allows us to clearly define localised occurrences. RQM's information foundation also strongly supports basing our understanding of metaphysical indeterminacy - and the open future - on the information available from different perspectives. The indeterminacy from one perspective, even if there is determinacy from another perspective, has real physical consequences in the form

---

[14] She also labels this a case of strong emergence (with no supervenience on the microlevel at all) - but it is not clear that this emergence is distinctly different from the way that standard quantum mechanics considers decoherence as a process of emergence (as considered in the previous chapter).

of interference effects. We cannot ignore this aspect of the theory in favour of the external eternalist view without losing significant explanatory value.

The issues surrounding sequences of outcomes over time - while not solved absolutely - give good reason to think that definiteness of past results should be based on robust physical records that persist into the future. At a microscopic, non-decoherent level definiteness is limited to a single event and both future and past (non-recorded) events should be treated as metaphysically indeterminate. This also suggests a way forwards for RQM in addressing the outstanding problems it has with cross-perspective links and the ontology of macroscopic, composite systems. Differentiating between microscopic correlations and robust macroscopic correlations produced by decoherence is a natural progression of RQM's claim that information is physical. The assumptions of continuity over time and a subjective experience comparable to our own should only apply to systems that form robust correlations. These are the correlations that can act as records to encode information over time that is accessible to observers.

As I have mentioned already, I do not intend for this view of time to rest on RQM alone. But RQM is a compelling example of how physics supports this model and conversely of how this model can help us interpret theories. This builds on the work of the last chapter on showing how records play an important role in connecting different levels of reality. The same basic ideas about how robust correlations come to exist through decoherence are used here; however, RQM does not need to accept quantum darwinism in its entirety. Instead, the same conclusions arise from applying only the same basic ideas about the emergence of robust correlations to an entirely different approach to quantum mechanics; this recurrence clearly demonstrates the broad applicability and widespread support for this model of time.

What is more, RQM is just one of a number of developing views that make use of perspectives and information. Other similar views include QBism, Perspectival Quantum Realism (Dieks 2022), Pragmatist Quantum Realism (Healey 2012), qubit based approaches (Höhn 2017), and several others. This family of information based, perspectival approaches to quantum mechanics offer a new way to think about how time appears in physics. Recent work explores the philosophy of time behind these views - for example Dieks (2022) presents his interpretation, perspectival quantum realism, and shows how it can be understood in terms of flow fragmentalism[15]. Many of them are also good candidates for localised temporal becoming. The idea of a quantum temporal becoming has been around for some time, although sometimes it has been judged negatively. For example, Callender (2017) argues that quantum temporal becoming is impossible because it forces us to accept an interpretation that has a preferred foliation of spacetime (at odds with relativity) and interpretations that allow this (e.g. non-relativistic Bohmian mechanics) seem to only provide a temporal becoming that is completely disconnected from experience. However, *localised* temporal becoming has no need for a preferred foliation and has found more success (for example Mariani & Torrengo (2021) apply localised temporal becoming to GRW as discussed in chapter 2). If we want to connect becoming to the experience of flow then information based, perspectival interpretations are even more promising, as the work of this chapter has shown. All these views aim to reconstruct quantum mechanics from simple principles about information. Due to this shared basis they all rely on the foundational importance of information and use perspectives to explain counterintuitive results of quantum mechanics such as interference and entanglement. Regarding the use of perspectives in QBism, Chris Fuchs has also recently said that

> "Quantum theory from the point of view of QBism tells us that reality itself is something more than any third-person perspective can capture. It is telling us that the universe is no block universe in the sense of that William James spells out: i.e., having "no ambiguous possibilities hidden in its womb," and where "the whole is in each and every part, and welds it with the rest into an absolute unity, an iron block."" (Crease & Sares 2021 p. 546).

This strongly aligns with the project of this thesis: to emphasise the ineliminable role of the internal perspective on time.

RQM, as a forerunner among these views, has been the best place to explore these

---

[15] Flow fragmentalism has many similarities to the model of temporal becoming presented here. See footnotes in chapter 2 for a brief comparison.

ideas in more precise detail. It also takes the most objectively realist stance among this family of views by rejecting that consciousness has any particular role (in contrast to QBism where the subjective experience of a human agent is central - QBism claims to be a realist interpretation but this is not obvious). But the lesson - that when interpreted from an internal perspective quantum mechanics supports the conclusion that the future is open and the past is fixed, due to the way that epistemic structures form in the world - may well be found in many other comparable interpretations. Combined with the work on quantum darwinism of the previous chapter - which is more naturally (though by no means absolutely) aligned with Everettian style interpretations - it is evident that an information based, emergent understanding of time is strongly supported by quantum mechanics, regardless of precisely how it is interpreted.

7 Conclusion

The goal of this thesis was to build a model of time's flow by following a programme of naturalised yet emergent and perspectival metaphysics. This rejects both the idea that flow is a psychological illusion and the claim that we must find flow in our most fundamental physics. Instead, I propose that objective flow can be found in physics by focusing on the theories that describe the realm of our experience. This means looking at how we interact with the world to gain information about the past and future (which is the basis for our experience of time), and identifying the physical structure that allows us to do this. It also means focusing on the macroscopic level of reality that we are placed in and the emergent theories that describe this. And going even further, it means acknowledging how our limited perspective - at a particular moment in time, a particular point in space, and with limited access to information - shapes the way that we can make inferences about other times. I will now summarise and bring together how these three threads - information, perspectives, and emergence - have been developed and how they work together to give new insights into the flow of time.

The account I give is this: flow is the localised becoming of events along a worldline that transforms an indeterminate future into a determinate past. Determinateness is rooted in the epistemic structure of the world; the future is open because it cannot be predicted with certainty from our embedded, localised, perspective on time, and the past is closed because physical records of it exist in the present that act as determinates for past claims.

I started with the idea that a model of the flow of time needs to grasp two things: a) how reality changes from moment to moment and b) what the transformative relationship between moments is. To achieve the first part of this I adopted localised temporal becoming but argued that it should be understood in strongly perspectival terms. If we place ourselves outside of and separate from time then we should be eternalist, but from the inside when we are embedded at a particular moment we should be presentists. In the embedded perspective, each event is present and exists from its own perspective while another event exists at the next moment, hence reality changes from moment to moment. The importance of the internal perspective led to the title of this thesis: learning about time from the *inside out*.

As for the the relationship between moments, this is where the open future comes in. Relative to a particular localised present moment, the future is indeterminate because there is no information existing in the present that can predict it (and hence settle the truth value of future claims). And past is fixed because of and to the extent that abundant records exist in the present that determine it. The transformation between them is what makes becoming a substantial account of flow. I argued in chapter 2 that this is compatible with the perspectival ontology of localised temporal becoming, but what remained to be shown was that this is not just important to our subjective experience or our intuitive way of thinking about time but also to the way that physics describes the objective world. This further link is what allows us to claim that flow is an essential and ineliminable part of reality that cannot be reduced to a psychological illusion. It takes us from a coherent metaphysical picture to one that we have good reason, motivated by naturalistic principles, to think that we *should* adopt.

To find this link I defended the view that the openness of the future is a perspectival metaphysical indeterminacy produced by the emergent epistemic structures in our world. I argued that we should take seriously the emergentist project and acknowledge that this plays an important role in our higher level theories.

Chapter 2 introduced the use of perspectives, and chapter 3 brought in emergence and information. But how are these related? The connection between perspectives and information is simple: our perspective dictates what information we have about the past and the future. We only ever directly experience the present but we make sophisticated inferences about other times by using the information available in the present, combined with our scientific theories, to predict and retrodict. Two further elements connect this to emergence: first is that our embedded perspective is situated at the macroscopic emergent level. Chapters 4-6 went on to show that the epistemic structure we typically use to make inferences also emerges and becomes robust at this level. Second, adding to this, our limitation to a *localised* present perspective is what makes emergent theories so effective at allowing us to make inferences about the past and the future. Emergent theories are particularly useful for dealing with incomplete, level-specific, and coarse-grained information. Likewise, instead of precisely retrodicting the past, we are forced to rely on macroscopic records - which are local systems - to make inferences (chapter 5 discussed how localisation leads to the essential properties of records, and the link between records and the screening off condition for emergence).

What is more, this aligns with the widely accepted reductionist project of time's arrow that looks at how all the temporal arrows can be reduced to the asymmetry of the thermodynamics, and at how this asymmetry emerges in statistical mechanics. Becoming - which we take to be the succession of moments in *one direction or another* - should likewise connect to emergence. To have a fundamentally asymmetric becoming would be at odds with the physics of asymmetry that is only found at higher levels of reality. The physics of time asymmetry also reinforces the connection between time and information. Explanations of thermodynamic asymmetry make frequent use of records and the associated accounts of the asymmetries of causation and agency reinforce this. It becomes increasingly apparent that the fact that we know about the past through records is far from being a subjective and trivial aspect of our experience but is instead indicative of how the world itself is distinctly asymmetric at the macroscopic level.

It is the fact that the limits on information are built into our theories that makes it clear that this is perspectival but *not* subjective. It is not only that we do not know the future but that we *cannot* know it; instead of arbitrary ignorance we have a well-defined but incomplete set of information that can be known. This gives a far broader physical basis for an open future than the literature on open futures currently considers (which is largely limited to the lightcone structure alone, if the physical basis is examined at all). And it much more clearly links the openness of the future to how we experience the world and the role that our epistemic inferences play in this.

The remaining three chapters laid out the physical justification for this model of time and showed how it involves the three elements of information, emergence, and perspectives.

I first focused on information and developed an account of records, defining them as correlations that are robust against noise. Records have received far less attention than the notion of probability, and I showed in chapters 4 and 5 that this account can be effectively applied in both thermodynamics and in quantum mechanics to provide new insights, such as how records feature in the debate about the thermodynamic cost of information, how records could connect to coarse-graining methods used to explain time asymmetry, and how records set a standard of classicality for quantum darwinism to aim for. This account also explains why records are predominantly macroscopic. Chapter 5 picked up this strand and gave a more detailed exploration of how records connect to accounts of emergence in the context of quantum decoherence.

Finally, I considered relational quantum mechanics as a case study for this model of time. The purpose of this was to illustrate that far from being a 'view from nowhere', physics can and frequently is interpreted in a perspectival way and that perspectives are closely linked to information. I showed that the model of time I have developed is a natural fit for this interpretation and that it suggests new ways to address current problems for the theory.

I have covered a broad range of different theories and interpretations and my intention is not to say that this model of time holds because *quantum darwinism or relational quantum mechanics specifically* are true. The aim of these chapters was, firstly, to explore the themes of emergence and perspectives along with their connection to information, and to demonstrate how this can be used to interpret physics. Secondly, these chapters show how quantum mechanics fits into this model of time: rather than being a unique and fundamental source of indeterminateness there is a strong line of argument, which comes from a range of different interpretations, that indicates that we can view quantum openness (and fixity) as information based and closely related to methods of inference and measurement. It also shows how the quantum world is continuous with higher level classical reality; this lends support to taking higher level indeterminateness seriously and to looking for the flow of time at an emergent level.

The clearest outcome of a perspectival, emergent, information based account is that there is no unique flow of time. Flow is localised along a worldline; this in itself has been accepted for some time since the results of special relativity have forced us to accept new ideas about time. But flow is also different at different levels of reality. The flow between microscopic events is distinct from the flow between macroscopic ones. But acknowledging that flow is a) the change in reality from moment to moment and b) the relation between moments which transforms indeterminateness to determinateness, naturally leads to this result. Reality is described differently at different levels and microscopic events are linked together by different laws than macroscopic ones.

The emergentist project tells us that we have higher level objects in our ontologies as well as fundamental ones. If we take this metaphysics seriously then the relativisation of flow to different levels is a necessary consequence. When we make claims about the future we formulate this in terms of the objects the world is populated with; but macroscopic objects do not reduce to clearly defined microscopic objects. This is what makes it meaningful to probabilistically predict the possible future macrostates of the world even when there is deterministic evolution of the microstate.

The difference between levels can also help us understand the debate between presentism and eternalism. If one looks at the type of examples that presentists often use to talk about time, it becomes clear that most of the reasoning used is about macroscopic objects of our direct experience. For example Prior (1968) talks about chairs, tables, people and horses. McTaggart (1908) uses the death of Queen Anne as an example of an event. And likewise the literature on the indeterminate truth values of future claims commonly uses macroscopic examples such as sea battles. Conceptually the motivation for prioritising the present is that *we* - complex macroscopic beings - have privileged experience of it that we do not have of the past and future. In contrast many of the arguments for eternalism come from looking at fundamental physics; for example Price (1996) looks at the microscopic fork asymmetry and denies that it can be found in collisions between particles. Although this is a generalisation and much of the reasoning can be imported to different scales, it reveals that the two views are to some extent focused on different levels. The perspectival relationship between presentism and eternalism can make sense of this. At lower levels of reality there is no significant epistemic structure in the world and there is little difference between the open future and the fixed past. At these levels the presentist perspective may not be of much value; the eternalist picture is perfectly adequate and potentially easier to work with. Temporal becoming still occurs between microscopic events; but the additional element of metaphysical asymmetry is weak or missing entirely. This element is what turned becoming into a substantial account of flow. However, when it comes to macroscopic reality, the eternalist picture is incomplete and the presentist perspective is needed to make sense of the open future. Here flow becomes an ineliminable aspect of the world.

What this thesis has achieved is to lay out the naturalistic basis for our everyday experience of time's flow. There are many questions which could lead on from this; for example one could look at the semantics of truthmakers and how this relates to metaphysical indeterminacy in more depth. I have relied on ideas about indeterminateness mainly from Wilson (2013), along with ideas from Barnes & Cameron (2010); further work could explore how a theory of truthmakers can be adapted to make sense of the level-dependent, perspectival claims I am relying on. But this is a secondary project, and what has been shown is that there is strong physical support for understanding time's flow in this way. And that information, perspectives and emergence provide a robust and valuable approach to answering these questions. To this extent the project of naturalistic metaphysics is satisfied; the features which must be included in our metaphysics are now clear. This is also not a comprehensive analysis of all the theories that use these ideas; further work could explore other areas using the ideas laid out here as a basis.

The results of this thesis can be summarised in three claims: 1) Time's flow is grounded on information: it can be identified by the way that information is gained from the world and used to predict or retrodict, and in how these concepts play a key role in physics. 2) It is perspectival: it is a result of the limitations that having a particular perspective puts on how you make inferences about the world. And 3) it is emergent: it is found in higher level theories describing the macroscopic realm of our experience.

References

Adlam, E. (2023). Do We Have any Viable Solution to the Measurement Problem?. *Foundations of Physics*, 53(2), 44.

Adlam, E. (2024). The Combination Problem for Relational Quantum Mechanics. *arXiv preprint arXiv:2401.15790*

Adlam, E., & Rovelli, C. (2023). Information is Physical: Cross-Perspective Links in Relational Quantum Mechanics. *Philosophy of Physics* 1(1): 4, 1–19.

Albert, D. (2000). Time and Chance.

Albert, D. (2015). The Difference between the Past and the Future. In *After Physics*, Harvard University Press.

Arthur, R. T. (2019). The Reality of Time Flow: Local Becoming in Modern Physics. Springer.

Bacciagaluppi, G. (2002). Remarks on space-time and locality in Everett's interpretation. In *Non-locality and Modality* (pp. 105-122). Dordrecht: Springer Netherlands.

Barad, K. (2007). *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning.* duke university Press.

Barnes, E. & Cameron, R. (2009) The Open Future: bivalence, determinism and ontology, *Philosophical Studies*, 146, 291-309

Barnes, E. & Cameron, R. (2011) Back to the Open Future. *Philosophical Perspectives*, 1-26

Barnes, E., & Williams, J. R. G. (2011). A theory of metaphysical indeterminacy.

Barnes, E. (2012). Emergence and fundamentality. *Mind*, *121*(484), 873-901.

Barnes, E. (2014). Fundamental indeterminacy. *Analytic Philosophy*, *55*(4), 339-362.

Belnap, N. (1992). Branching Space-time. *Synthese, 92*(3), 385-434.

Belnap, N. (2003). Branching Space-Time, Postprint. Retrieved from http://philsci-archive.pitt.edu/id/eprint/1003

Belnap, N. (2012). Newtonian Determinism to Branching Space-Times Indeterminism in Two Moves. *Synthese, 188*, 5-21.

Bigelow, J. (1996). Presentism and Properties. In J. Tomberlin (Ed.), *Philosophical Perspectives* (pp. 35-52). Cambridge, MA: Blackwell

Born, M., & Jordan, P. (1925). Zur Quantenmechanik. *Z. Phys.*, 34, 854-888. doi:10.1007/BF013 28531

Born, M., Jordan, P., & Heisenberg, W. (1926). Zur Quantenmechanik II. *Z. Phys.*, 34, 557-615. doi:10.1007/BF01379806

Brading, K. (2013). Presentism as an empirical hypothesis. *Philosophy of Science*, *80*(5), 1101-1111.

Broad, C. D. (1923) Scientific Thought London: Routledge & Kegan Paul

Brown, M. J. (2009). Relational Quantum Mechanics and the determinacy problem. *The British Journal for the Philosophy of Science*, 60(4), 679-695.

Callender, C. (2017). What makes time special?, Oxford University Press.

Calosi, C., & Mariani, C. (2020). Quantum Relational Indeterminacy. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 71, 158-169.

Cameron, R. (2015). The moving spotlight. An essay on time and ontology. Oxford: Oxford University Press

Candiotto, L. (2017). The Reality of Relations. *Giornale di Metafisica*, 2, 357-551.

Cavalcanti, E. G., Di Biagio, A., & Rovelli, C. (2023). On the consistency of relative facts. arXiv preprint arXiv:2305.07343.

Chang, H. (2022). *Realism for realistic people.* Cambridge University Press.

Crease, R. P., & Sares, J. (2021). Interview with physicist Christopher Fuchs. *Continental Philosophy Review*, 1-21. pp.546.

Crisp, T. (2007). Presentism and the Grounding Objection. *Nous*, 41(1), 90-109

Cuffaro, M. E., & Hartmann, S. (2021). The open systems view. *arXiv preprint arXiv:2112.11095.*

Deng, N. (2013a). Our experience of passage on the B-theory. *Erkenntnis*, *78*, 713-726.

Deng, N. (2013b). On explaining why time seems to pass. *The Southern Journal of Philosophy*, *51*(3), 367-382.

Deng, N. (2017). Making Sense of the Growing Block View. *Philosophia, 45*(3), 1113-1127

Deng, N. (2023). Time, Grounding, and Esoteric Metaphysics. *The Monist*, *106*(3), 287-300.

Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27-51.

Di Biagio, A., & Rovelli, C. (2021). Stable facts, relative facts. *Foundations of Physics*, 51(1), 1-13.

Di Biagio, A., Donà, P., & Rovelli, C. (2021). The arrow of time in operational formulations of quantum theory. *Quantum*, *5*, 520

Dieks, D. (2006). Becoming, Relativity and Locality. In D. Dieks, *The Ontology of Spacetime.* Elsevier.

Dieks, D. (2022). Perspectival quantum realism. *Foundations of Physics*, *52*(4), 95

Dirac, P. (1956). The Fundamental Equations of Quantum Mechanics. *Proc. R. soc. Lond. Ser. A, 109*, 645-653. doi:10.1098/rspa.1925.0150

Dorato, M. (2000). Becoming and the Arrow of Causation. *Philosophy of Science*, *67*(S3), S523-S534

Dorato, M. (2006). Absolute becoming, relational becoming and the arrow of time: Some non-conventional remarks on the relationship between physics and metaphysics. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *37*(3), 559-576.

Dorato, M. (2016). Rovelli's relational quantum mechanics, monism and quantum becoming. In A. Marmodoro, & A. Yates (Eds.), *The Metaphysics of Relations* (pp. 290-324). Oxford, UK: Oxford University Press. Retrieved from http://arxiv.org/abs/1309.0132

Dowker, F., & Kent, A. (1996). On the consistent histories approach to quantum mechanics. *Journal of Statistical Physics*, *82*, 1575-1646.

Dummett, M. (1954). Can an Effect Precede Its Cause? *Proceedings of the Aristotelean Society, Supplementary Volume, 38*, 27-44.

Dummett, M. (1964). Bringing about the Past. *Philosophical Review, 73*, 338-359.

Durt, T. (2010). Anthropomorphic quantum Darwinism as an explanation for classicality. *Foundations of Science*, *15*, 177-197

Earman, J. (2008). Reassessing the Prospects for a Growing Block Model of the Universe. *International Studies in the Philosophy of Science, 22*(2), 135-164.

Einstein, A., Tolman, R. C., & Podolsky, B. (1931). Knowledge of past and future in quantum mechanics. *Physical Review*, *37*(6), 780.

Elga, A. (2005). Isolation and Folk Physics. In H. Price, & R. Corry (Eds.), *Russell's Republic: The Place of Causation in the Constitution of Reality.* Oxford University Press.

Everett, H. (1957). Relative State Formulation of Quantum Mechanics. *Rev. Mod. Phys. 29*, 454-462.

Fernandes, A. (2023). Time, Flies, and Why We Can't Control the Past. In B. Loewer, B. Weslake, & E. Winsberg (Eds.), *The Probability Map of the Universe: Essays on David Albert's Time and Chance*, Harvard University Press. (pp. 312–334).

Fields, C. (2010). Quantum Darwinism requires an extra-theoretical assumption of encoding redundancy. *International Journal of Theoretical Physics*, *49*(10), 2523-2527.

Fields, C. (2011). Classical system boundaries cannot be determined within quantum Darwinism. *Physics Essays, 24*(4), 518.

Fine, K. (2005). Tense and Reality. In *Modality and Tense: Philosophical Papers.* Oxford University Press.

Franklin, A. (2022). Incoherent? No, Just Decoherent: How Quantum Many Worlds Emerge. *Philosophy of Science*, 1-22.

Franklin, A., & Robertson, K. (2022). Emerging into the Rainforest: Emergence and Special Science Ontology. http://philsci-archive.pitt.edu/19912

Frisch, M. (2023). Causes, randomness, and the past hypothesis. In Loewer, B., Weslake, B., & Winsberg, E. (Eds.). *The Probability Map of the Universe: Essays on David Albert's Time and Chance.* Harvard University Press.

Gell-Mann, M., & Hartle, J. B. (1993). Classical equations for quantum systems. *Physical Review D*, 47(8), 3345

Gell-Mann, M. & Hartle, J.B. (1994), "Equivalent Sets of Histories and Multiple Quasiclassical Domains", Preprint UCSBTH-94-09, gr-qc/9404013, submitted to gr-qc 8 April 1994

Gell-Mann, M., & Hartle, J. B. (1995). Strong decoherence. *arXiv preprint gr-qc/9509054.*

Grandjean, V. (2022). *The Asymmetric Nature of Time: Accounting for the Open Future and the Fixed Past* (p. 216). Springer Nature.

Greaves, H., & Wallace, D. (2014). Empirical consequences of symmetries. *The British Journal for the Philosophy of Science.*

Halliwell, J. J. (1999). Somewhere in the universe: Where is the information stored when histories decohere?. *Physical Review D*, 60(10), 105031.

Hartle, J. (2004). The Physics of 'Now'. *American Journal of Physics*, 73(2), 101-109.

Hawking, S. (1994). The No Boundary Condition and the Arrow of Time. In J. J. Halliwell, J. P'erez-Mercader, & W. Zurek (Eds.), *Physical Origins of Time Asymmetry.* Cambridge: Cambridge University Press.

Healey, R. (2012). Quantum theory: A pragmatist approach. *The British Journal for the Philosophy of Science.*

Heisenberg, W. (1925). Uber quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen. *Z. Phys*, 33, 879-893.

Hemmo, M., & Shenker, O. R. (2012). *The Road to Maxwell's Demon: Conceptual Foundations of Statistical Mechanics.* Cambridge University Press.

Höhn, P. A. (2017). Quantum theory from rules on information acquisition. *Entropy, 19*(3), 98.

Hernández-Orozco, S., Hernández-Quiroz, F., & Zenil, H. (2018). Undecidability and irreducibility conditions for open-ended evolution and emergence. *Artificial Life, 24*(1), 56-70.

Horwich, P. (1988). *Asymmetries in Time: Problems in the Philosophy of Science.* 2nd ed. Cambridge, MA: MIT Press.

Huggett, N. (2023). Reading the Past in the Present, In B. Loewer, B. Weslake, E. Winsberg (Eds), *The Probability Map of the Universe: Essays on David Albert's Time and Chance*

Ismael, J. (2006). Saving the baby: Dennett on autobiography, agency, and the self. *Philosophical Psychology*, 19(3), 345-360.

Ismael, J. (2009). Probability in deterministic physics. *The Journal of Philosophy, 106*(2), 89-108.

Ismael, J. (2016a). From Physical Time to Human Time. In R. M. Dolev Y. (Ed.), *Cosmological and Psychological Time* (Vol. vol 285): Springer

Ismael, J. (2016b). *How physics makes us free.* Oxford University Press.

Ismael, J. (2023). The Open Universe: Totality, Self-reference and Time. *Australasian Philosophical Review*, 1-16.

James, W. (1890). *The Principles of Psychology.* New York: Henry Holt and Company.

Joos, E. (2000). Elements of environmental decoherence. *Decoherence: Theoretical, experimental, and conceptual problems*, 1-17.

Joos, E., & Zeh, H. D. (1985). The emergence of classical properties through interaction with the environment. *Zeitschrift für Physik B Condensed Matter*, 59(2), 223-243.

Kierland, B. (2013). Grounding Past Truths: Overcoming the Challenge. In R. Ciuni, K. Miller, & G. Torrengo (Eds) *New Papers on the Present: Focus on Presentism*, Philosophia

Korbicz, J. K. (2021). Roads to objectivity: quantum darwinism, spectrum broadcast structures, and strong quantum darwinism–a review. *Quantum*, 5, 571.

Ladyman, J., & Robertson, K. (2013). Landauer defended: reply to Norton. *Studies in the History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physic,* 44(3), 263-271.

Ladyman, J., & Robertson, K. (2014). Going round in circles: Landauer vs Norton on the thermodynamics of computation. *Entropy*, 16(4), 2278-2290.

Ladyman, J., Presnell, S., & Short, A. J. (2008). The use of the information theoretic entropy in thermodynamics. *Studies in the History and Philosophy of Modern Physics,* 39, 315-324.

Ladyman, J., Presnell, S., Short, A. J., & Groisman, B. (2007). The connection between logical and thermodynamic irreversibility. *Studies in the History and Philosophy of Modern Physics*, 38, 58-79.

Ladyman, J. & Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized.* Oxford University Press. With David Spurrett and John Collier

Landauer, R. (1961). Irreversibility and heat generation in the computing process, *IBM Journal of Research and Development*, 5, pp. 183–191

Laudisa, F., & Rovelli, C. (2019). Relational Quantum Mechanics. *Retrieved from Stanford Encyclopedia of Philosophy*

Leininger, L. (2021a). Temporal B-coming: Passage without presentness. *Aus-*

tralasian Journal of Philosophy, 99(1), 130-147.

Leininger, L. (2021b). Coordination and Coming to Be. Australasian *Journal of Philosophy* 101 (1):213-227.

Lewis, D. (1979). Counterfactual Dependence and Time's Arrow. *Nous, Special Issue on Counterfactuals and Laws*, 13(4), 455-476

Lewis, D. (1986). On the Plurality of Worlds. Oxford: Blackwell.

Lipman, M. A. (2015). On Fine's fragmentalism. *Philosophical Studies*, *172*, 3119-3133.

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

Loewer, B. (2023) The Mentaculus: A Probability Map of the Universe, In Loewer, B., Weslake, B., & Winsberg, E. (Eds.). *The Probability Map of the Universe: Essays on David Albert's Time and Chance.* Harvard University Press.

Lowe, E. J. (1987). The indexical fallacy in McTaggart's proof of the unreality of time. *Mind*, *96*(381), 62-70

MacFarlane, J. (2003). Future Contingents and Relative Truth. *The Philosophical Quarterly, 53*, 321-336.

Mariani, C., & Torrengo, G. (2021a). The indeterminate present and the open future. *Synthese*, *199*(1-2), 3923-3944.

Mariani, C., & Torrengo, G. (2021b). The Metaphysics of Passage in Dynamical Reduction Models of Quantum Mechanics. In *Ockhamism and Philosophy of Time: Semantic and Metaphysical Issues Concerning Future Contingents* (pp. 147-172). Cham: Springer International Publishing.

Mariani, C. (2022). Indeterminacy: Deep but not rock bottom. *Analytic Philosophy*, *63*(1), 62-71.

Markosian, N. (1995). The open past. *Philosophical studies: an international journal for philosophy in the analytic tradition*, *79*(1), 95-105.

Markosian, N. (2013). The truth about the past and the future. *Around the Tree: Semantic and Metaphysical Issues Concerning Branching and the Open Future*, 127-141.

Maroney, O. J. (2010). Does a Computer have an Arrow of Time? *Foundations of Physics*, 40(2), 205-238.

Maroney, O. J., & Timpson, C. G. (2018). How is there a physics of information? On characterizing physical evolution as information processing. *Physical perspectives on computation, computational perspectives on physics*, 103-126.

Martin-Dussaud, P., Rovelli, C., & Zalamea, F. (2019). The notion of locality in relation quantum mechanics. *Foundations of Physics*, 49(2), 96-106.

Massimi, M. (2022). Perspectival realism. Oxford University Press.

Maudlin, T. (2007). *The metaphysics within physics.* Oxford University Press.

Maxwell, N. (1985). Are Probabilism and Special Relativity Compatible? *Philosophy of Science, 52*, 23-43.

McCall, S. (1994). A Model of the Universe. Oxford: Oxford University Press.

McKenzie, K. (2011). Arguing against fundamentality. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics, 42*(4), 244-255.

McTaggart, J. (1908). The Unreality of Time. *Mind, 17*, p. 457-474.

Mozersky, M. J. (2015). *Time, Language, and Ontology: The World From the B-Theoretic Perspective*: Oup Oxford.

Muciño, R., E. Okon, and D. Sudarsky. 2021. "Assessing Relational Quantum Mechanics." Synthese 200, no. 5: 1–26.

Myrvold, W. C. (2021a). *Beyond chance and credence: A theory of hybrid probabilities.* Oxford University Press, USA.

Myrvold, W. C. (2021b). Shakin' all over: Proving Landauer's principle without neglect of fluctuations. Forthcoming in BJPS; online at https://www.journals.uchicago.edu/doi/10.108

Myrvold, W. C., & Norton, J. D. (2023). On Norton's "... Shook..." and Myrvold's "Shakin'...".

Nielson, M. A., & Chuang, I. L. (2000). *Quantum Computation and Quantum Information.* Cambridge: Cambridge University Press.

Norton, J. D. (2011). Waiting for Landauer. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 42(3), 184-198.

Norton, J. D. (2013a). The end of thermodynamics of computation: A no-go result. *Philosophy of Science*, 80(5), 1182-1192.

Norton, J. D. (2017b). The worst thought experiment. In *The Routledge companion to thought experiments* (pp. 454-468). Routledge.

Norton, J.D. (2017a) Thermodynamically reversible processes in statistical physics. *American Journal of Physics*, *85*(2), 135-145.

Oaklander, L. N. (2015). Temporal Phenomena, Ontology and the R-theory. *Metaphysica, 16*(2), 253–269

Okon, E., & Sudarsky, D. (2016). Less decoherence and more coherence in quantum gravity, inflationary cosmology and elsewhere. *Foundations of Physics*, *46*, 852-879.

Pienaar, J. (2021) A quintet of quandaries: Five no-go theorems for relational quantum mechanics. *Foundations of Physic*s, 51(5)

Placek, T., & Belnap, N. (2012). Indeterminism is a modal notion: branching spacetimes and Earman's pruning. *Synthese, 187*, 441-469.

Pooley, O. (2013). Relativity, the Open Future, and the Passage of Time. *Proceedings of the Aristotelian Society, 113*(3/3), 321-363

Popescu, S., Short, A.J, & Winter, A., (2005). The Foundations of Statistical Mechanics from Entanglement: Individual States vs. Averages. *Preprint, arXiv:quant-ph/0511225.*

Price, H. (1996). *Time's Arrow and Archimedes' Point.* New York: Oxford University Press.

Price, H. (2007). Causal perspectivalism. In H. Price & R. Corry (Eds.), Causation, physics and the constitution of reality: Russell's Republic revisited. Oxford: Oxford University Press

Prior, A. (1968). Changes in Events and Changes in Things. In *Papers on Time and Tense* (p. 1-14). Oxford: Oxford University Press.

Prokopenko, M., Harré, M., Lizier, J., Boschetti, F., Peppas, P., & Kauffman, S. (2019). Self-referential basis of undecidable dynamics: From the Liar paradox and the halting problem to the edge of chaos. *Physics of life reviews*, *31*, 134-156.

Putnam, H. (1967). Time and Physical Geometry. *The Journal of Philosophy, 64*(8), 240-247.

Reichenbach, H. (1956). *The Direction of Time.* (M. Reichenbach, Ed.) Berkeley and Los Angeles: University of California Press.

Riedel, C. J., & Zurek, W. H. (2010). Quantum Darwinism in an everyday environment: Huge redundancy in scattered photons. *Physical review letters*, 105(2), 020404.

Riedel, C. J., Zurek, W. H., & Zwolak, M. (2016). Objective past of a quantum universe: Redundant records of consistent histories. *Physical Review A*, *93*(3), 032126.

Rietdijk, C. (1966). A Rigorous Proof of Determinism Derived from the Special Theory of Relativity. *Philosophy of Science, 33*(4), 341-344.

Robb, A. A. (1914). *A Theory of Time and Space.* Cambridge: The University Press.

Robertson, K. (2020). Asymmetry, Abstraction, and autonomy: Justifying coarse-graining in statistical mechanics. *The British Journal for the Philosophy of Science*, 71(2), 547-579.

Robertson, K. (2021). Autonomy generalised; or, Why doesn't physics matter more?. *Ergo.*

Robertson, K., & Prunkl, C. (2023) Is thermodynamics subjective?. *Philosophy of*

Science, 1-16.

Rocha, G. R., Rickles, D., & Boge, F. J. (2021). A Brief Historical Perspective on the Consistent Histories Interpretation of Quantum Mechanics. *arXiv preprint arXiv:2103.05280.*

Ross, D. (2000). "Rainforest realism: A Dennettian theory of existence" . In D. Ross, A. Brook, D. Thompson (Eds) *Dennett's Philosophy: A Comprehensive Assessment*, MIT Press.

Rovelli, C. (1996). Relational Quantum Mechanics. *International Journal of Theoretical Physics*, 35(8), 1637-1678.

Rovelli, C. (2018). Space is blue and birds fly through it. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2123).

Rovelli, C. (2019). Neither Presentism nor Eternalism. *Foundations of Physics, 49*(12), 1325-1335.

Savitt, S. (2002). On Absolute Becoming and the Myth of Passage. In C. Callender (Ed.), *Time, Reality & Experience* (Vol. 50, p. 153-).

Savitt, S. (2006). Presentism and eternalism in perspective. *Philosophy and Foundations of Physics*, 1, 111-127.

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

Schlosshauer, M. (2019). Quantum decoherence. *Physics Reports*, 831, 1-57.

Schulman, L. (2005). A Computer's Arrow of Time. *Entropy*, 7(4), 221-233.

Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A*, 52, 2493-2496.

Sider, T. (2011). Writing the Book of the World. Oxford: Oxford University Press.

Skow, B. (2009). Relativity and the Moving Spotlight. *Journal of Philosophy, 106*, 666-678.

Stein, H. (1968). On Einstein-Minkowski Space-Time. *Journal of Philosophy, 65*, 5-23.

Stradis, A. (2021). Memory, the fork asymmetry, and the initial state. *Synthese*, 1-25.

Strevens, M. (2011). Probability out of determinism. *Probabilities in physics*, *146*.

Szilard, L. (1929). "Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen," *Zeitschrift für Physik*, 53(1929), pp. 840-56

Timpson, C. G. (2013). *Quantum information theory and the foundations of quantum mechanics.* OUP Oxford.

Thomas, E. (2019). The roots of CD Broad's growing block theory of time. *Mind*, *128*(510), 527-549.

Torre, S. (2011). The Open Future. *Philosophy Compass*, 6, 360-373

Torrengo, G., & Iaquinto, S. (2019). Flow fragmentalism. *Theoria*, *85*(3), 185-201.

Torrengo, G., & Iaquinto, S. (2020). The invisible thin red line. *Pacific Philosophical Quarterly*, *101*(3), 354-382.

Uffink, J. (2002). "Essay review: Time and Chance", Studies in History and Philosophy of Modern Physics 33, 555-563.

Van Fraassen, B. (2010). Rovelli's World. *Foundations of Physics*, 40(4), 390-417.

Wallace, D. (2012). *The emergent multiverse: Quantum theory according to the Everett interpretation.* Oxford University Press.

Wallace, D. (2017). The Nature of the Past Hypothesis. In K. Chamcham, J. Silk, J. D. Barrow, & S. Saunders (Eds.), *The Philosophy of Cosmology* (pp. 486-499). Cambridge University Press.

Wallace, D. (2019a). Who's afraid of coordinate systems? An essay on representation of spacetime structure. *Studies in history and philosophy of science part B: Studies in history and philosophy of modern physics*, *67*, 125-136.

Wallace, D. (2019b). Observability, redundancy and modality for dynamical symmetry transformations. http://philsci-archive.pitt.edu/16622/

Wallace, D. (2020). Fundamental and emergent geometry in Newtonian physics. *The British Journal for the Philosophy of Science.*

Wang, Y. (2012). Quantum Computation and Quantum Information. *Statistical Science*, 27(3), 373-394.

Weatherall, J. O. (2016). Maxwell-Huygens, Newton-Cartan, and Saunders-Knox space-times. *Philosophy of Science*, *83*(1), 82-92.

Wilson, J. (2013). A determinable-based account of metaphysical indeterminacy. *Inquiry*, 56(4), 359-385.

Wilson, J. (2016). Are there indeterminate states of affairs? Yes. In *Current controversies in metaphysics* (pp. 105-119). Routledge.

Wolpert, D. H. (1992). Memory Systems, Computation, and the Second Law of Thermodynamics. *International Journal of Theoretical Physics*, 31, 743-785.

Woodward, J. (2021). Explanatory autonomy: the role of proportionality, stability, and conditional irrelevance, *Synthese* 198.1, 237–265

Zeh, H. (1993). There are no quantum jumps, nor are there particles! *Phys. Lett. A*, 172, 189-192.

Zwolak, M., & Zurek, W. H. (2017). Redundancy of einselected information in quantum Darwinism: The irrelevance of irrelevant environment bits. *Physical Review A*, *95*(3), 030101.

Zurek, W. (1994). Preferred Sets of States, Predictability, Classicality, and Environment-Induced Decoherence. In J.J. Halliwell, J. Pérez-Mercader, & W.H. Zurek (Eds.), *Physical Origins of Time Asymmetry* . Cambridge: Cambridge University Press. 175-212

Zurek, W. (2006). Decoherence and the transition from quantum to classical - revisited. *Quantum Decoherence*, 1-31.

Zurek, W. (1982). Environment-induced superselection rules. *Physical review D*, 26(8), 1862.

Zurek, W. (1998). Decoherence, einselection and the existential interpretation (the rough guide). *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 356(1743), 1793-1821.

Zurek, W. (2009). Quantum darwinism. *Nature physics*, 5(3), 181-188.

Zurek, W. (2022). Quantum Theory of the Classical: Einselection, Envariance, Quantum Darwinism and Extantons. *Entropy*, 24(11), 1520.