Noisy Records

Abstract

Records of the past are a pervasive feature of our world that are used in many areas of both physics and philosophy, but out best account of records is overidealised and leaves many questions unanswered. I will present a new account of records in terms of robustness against noise. This account highlights previously overlooked features of records: their use of redundancy and their reliance on techniques common in the physics of macroscopic phenomena, such as coarse-graining. These features have implications for how we explain the record asymmetry through statistical mechanics as well as for related debates such as that surrounding Maxwell's demon.

Keywords: records; time asymmetry; statistical mechanics; thermodynamics; Maxwell's demon

1 Introduction

Records of the past are a prevalent and inescapable part of our world. They form a huge element of our everyday experience of temporal asymmetry. They also feature heavily in discussions of time asymmetry in statistical mechanics as well as other related debates. The most successful analysis of records to date is given by Albert (2000; 2015) as part of his wider analysis of time asymmetry in statistical mechanics. Albert's programme aims at providing an overarching explanation for all temporal asymmetries in the world and argues that they reduce ultimately to statistical mechanics; records are a key part of this that ground many other arrows such as the causal asymmetry and the asymmetry of agency. However, his account of records is limited by overidealisation which, while understandable within the overall project he is aiming for, leaves a number of questions unanswered and overlooks crucial aspects of what a record is.

I will give a new account of records and how they fit into statistical mechanics, building on Albert's basic conception. 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 are *robust*, where robustness is understood in terms of robustness against noise.[1] My aim is to show that noise is not just a practical afterthought to how records work, but a central part of their definition that brings to light features that are essential to explaining how records function and how they relate to temporal asymmetry. Namely that records require the use of *redundancy* and that modelling records uses methods such as coarse-graining and statistical averaging.

The scope of this account is in one sense broad; it aims - as Albert's account does - to provide a general template for all the different sorts of records we might find in the world. However, it is also directed specifically at how records relate to statistical mechanics and the asymmetry therein, and as such the analysis will be geared towards that goal. Many records in the world cannot be easily characterised by the machinery of statistical mechanics and will be harder to fit into this account (even if we accept the idea that ultimately all asymmetry in the world comes down to statistical mechanics the connections may be complex). But through examples I will endeavour to show how the account applies more generally, with the caveat that each individual type of record will present its own complications that a general account can never completely address.

[1]Of course no account is completely without idealisation, and in this paper I will discuss many methods, such as coarse-graining and statistical methods, that 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 external perturbations 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.

Alongside presenting this account, I will explore two main philosophical implications that come out of it. These implications will demonstrate why it is essential to consider these previously overlooked features of records and how they can help our discussions of asymmetry in statistical mechanics.

The first implication is that this account of records provides a simple diagnosis as to why Maxwell's demon is impossible and what is going on in the literature surrounding this. Maxwell's demon is a thought experiment about how an intelligent agent could manipulate the fluctuations of individual particles in a system to manifest a violation of the second law of thermodynamics. This requires that the agent has a memory device, in particular one that functions using thermodynamically reversible processes on single particle systems. It has become common to refute Maxwell's demon by appealing to Landauer's principle (although there is significant controversy around proving this principle), which states that there is a thermodynamic cost to information processing that will balance out any violation (e.g. see Robertson (2021)). The account of records given here, however, provides a simpler explanation. Maxwell's demon fails regardless of this principle because the robustness - and use of redundancy - that records require cannot be achieved using thermodynamically reversible processes on single particle systems. As a result, it becomes clear that it is impossible to have an intelligent agent that is able to manipulate microscopic fluctuations into a violation of the second law of thermodynamics. This supports and strengthens the conclusions of Norton (2013), who argues that noise disrupts any microscopic computation in the reversible regime, and that this is what ultimately overcomes the demon rather than Landauer's principle; but it shows that this is not just a practical failure but comes from the definition of records.

The second implication is that identifying how records are connected to methods such as coarse-graining gives new ways to explain the record asymmetry. It indicates that we should move away from Albert's account, which focuses primarily on establishing a basic link between records and the past hypothesis through ready states (an essential feature of his account that I will argue plays no actual role in how we use records), and instead focus on the methods of coarse-graining and statistical averaging that records require. Wallace (2013) provides an alternate programme to Albert's for explaining macroscopic asymmetry that focuses on exactly these sorts of methods; he argues that macro-dynamics in physics emerge from micro-dynamics via these methods. Their effectiveness, and asymmetry, is grounded on an alternate version of the past hypothesis. Wallace's programme currently gives no account of records; one of his programme's drawbacks is that, while it works well within physics, it does not show how asymmetry in statistical mechanics is a basis for many of the other temporal asymmetries in the way that Albert's programme does. I will show how records are connected to Wallace's arguments, and how this could be a important avenue for future exploration into why we have a distinctive record asymmetry.

I will proceed by first, in section 2, presenting Albert's account of records and its drawbacks, then presenting my own account in section 3. My account looks at how noise contributes to the definition of what a record is as well as the more general role of stability in ensuring robustness against noise. I will then examine in section 4 a specific way to ensure robustness against noise: redundancy. First, I will look at it in the practical context of error correction coding in computing, and then I will show the theoretical importance of this idea by demonstrating how it helps us understand discussions around Maxwell's demon. Finally, I end in section 5 with a discussion of how records make use of methods such as coarse graining and the philosophical implications of this for explaining temporal asymmetry through Wallace's programme.

2 Records as Measurement: Albert's Ready State Account

While various theories to explain records have been proposed, Albert's (2000; 2015) statistical mechanical account is by far the most detailed and successful and it is the only one I will explore in depth.[2] 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.

Albert sees inference by measurement as a simple process in which we have information about an initial state and a final state (the record) and use this to make inferences about what happened in the intervening period. In more detail: a measuring 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 which uses only the current state and the dynamics.) We use the ready state to calculate what the evolution of the system would be if it remained isolated and we can work out what the expected current state should be. Any deviations between this expected current state and the actual current state indicates that isolation has been broken and an interaction with another system has occurred in the interval between

[2]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. Another account is the fork asymmetry account (Horwich 1988; Stradis 2021).

the ready state and the current time. The differences between the actual state and the expected state give information about this interaction. Albert uses a simple scenario to exemplify the use of ready states: a ready state shows that a 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. We can also make certain calculations about how this other ball must have been moving. This basic model, Albert claims, is the basis of all types of records, ranging from formal scientific measurements to everyday phenomena such as photographs and footprints (of course how well the model of fits varies and some types of records may have unique factors to consider - but the account is supposed to be a general description of the process).

Albert argues that the use of a ready state is what leads to the record asymmetry. To have knowledge of the ready state we require a second, independent record of it. This record must be modelled in the same way and will in turn require a third record of *its* ready state. This leads to an infinite regress that can only be stopped by the past hypothesis - a condition on the initial state on the universe that acts as the "mother (as it were) of all ready conditions" (Albert, 2015). Albert's account of asymmetry within statistical mechanics also comes down to the past hypothesis, and hence this connection between records and the past hypothesis links records in his overall programme for explaining all temporal asymmetries through statistical mechanics.

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 leading up to 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.[3]

Albert's account assumes a simple isolated interaction between just two systems. In reality there will be many systems interacting, each changing the state of the record 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 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 systems deliberately designed

[3] In the context of Albert's work this is understandable. His focus is not on the everyday use of records but on proving records reliable in the face of the *reversibility problem* in statistical mechanics that implies that we should retrodict a higher entropy past despite our records to the contrary. The details of inferences are left as what Albert calls a "foggy sense" of the ready states. Huggett (2023) raises similar concerns and goes into more detail about how Albert's account relates to local inferences, also considering the role of noise. He emphasises that the past hypothesis is a precondition to records and their asymmetry but not a direct part of our reasoning.

as measurement devices try to achieve as high a level of isolation as possible (consider the systems at CERN for instance) – 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 expected evolution based on the ready state 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 expected current state *should* be. Without a precise expected state we cannot identify how, or even if, an interaction has occurred to change the state of the record.

While, in these cases, 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 animals passing and is identifiable (at least to an expert tracker) despite there being no mechanisms shielding it. The cosmic microwave background radiation is a reliable record of the big bang, and many more. 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.

3 Non-Idealised Records

I will now present a new account of records in terms of their robustness against noise. Working noise into our definition of records, rather than as an operational afterthought, helps us identify key features of records that have been previously overlooked but are essential for understanding records.

What Albert's account successfully identified was that records are things in the world that allow us to make direct inferences to some past fact (in his account the inference is to the interaction that altered the expected state of the record into the actual state), as opposed to retrodiction which evolves the present state backwards to an arbitrary time. Retrodiction can be used to tell us about any time we choose, records link back to something specific. Additionally records tell us about other systems that the record system has interacted with, while retrodiction tells us about the past state of just one system (the system whose present state we evolve backwards). I will call the connection a record has to the past a direct *correlation* to some aspect of the past. This correlation can be a probabilistic connection (i.e. the record raises the probability of the past fact being true[4]), a causal link (i.e. the record is caused by this aspect of the past - although this may be circular if we then want records to ground our account of causation such as in Loewer (2023)), or any other sort of link that grounds the inference we make between the record and the past.[5] In essence the idea is that if this aspect of the past had been different, the record would be different too. (This aspect of the past may be a specific event such as the collision between billiard balls, some feature of a system such as the shape of a fossil, or a continuous past fact such as the atmospheric composition over time that is captured by an ice core sample.)

What makes a record informative is when this correlation is clearly identifiable and we can reliably establish this connection to some aspect of the past. This allows us to skip the more laborious retrodictive task of evolving the entire present state backwards and instead make direct inferences about the past. For this to be possible, the connection between the record and the past must stand out against the many different interactions and influences the the record system will inevitably be exposed to. That specific aspect of the past must be the deciding fact of what the current state the record is, while all other interactions produce only negligible changes. Taking Albert's billiard ball example: for us to identify the link between the current stationary state of the ball (the record) and the collision in the past we must assume that air resistance, friction, other glancing collisions that may have occurred, etc all had negligible impact on the balls motion. In this case these factors are easy to ignore, but in many cases they will not be and singling out the connection will be a much greater challenge. What this comes down to is that the record - and its connection to the past - must be *robust*. This means both that the connection is clearly identifiable in the first place and that the record system retains its connection to the past over time (although some records may be fleeting). I argue that this robustness is not an afterthought to improve functionality but a key theoretical consideration.

3.1 Robustness Against Noise

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 modifica-

[4]This is probably closest to what Albert's account proposes and how he talks about correlations. It is also similar to how records are thought of in Stradis (2021), which connects Albert's account to the fork asymmetry.

[5] We can model this inference using something like Norton's (2021) material theory of induction, which is a good fit for the account of records proposed here with its emphasis on contextuality and background information.

tions to a useful signal; it obscures the information in the signal and makes it harder to read. If we think of records as encoding 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 and disrupt the connection between the record and the specific aspect of the past that it is recording.

Unwanted modifications to a *desired* signal makes noise a very anthropocentric concept. 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 the term noise in Dennett (1991) where he considers the existence of real and objective patterns arising in nature, which 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 and records as well. According to Dennett, a real pattern is identified when it is simpler to describe the state of the world in terms of a pattern with a certain amount of noise affecting it than it is to give information about every detail of the state individually.[6] The patterns are identifiable despite noise partially obscuring them, 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. In signalling, the existence of a well-defined signal is presupposed; in Dennett's usage however, making the distinction between pattern and noise is what defines the pattern. Records can be thought of in this sense: by making the distinction between the record - which is informatively correlated to some aspect of the past - and noise, we are defining what the record is. Dennett's idea of patterns has been used in the emergence lit-

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

erature to capture how we can have emergent higher level objects, for example tables and chairs (Wallace 2013; 2024). These objects are stable patterns in the underlying molecular description, and it is usually far more efficient to talk in terms of these patterns than in terms of molecular physics. What I am proposing here for records is similar, although less general. Records are specifically robust patterns that come with a correlation to the past and hence encode information in a certain way. Take the cosmic microwave background radiation (CMBR). This is not itself an object, it is a collection of photons, but when we identify these photons as linking back to the big bang it becomes a record. To do this we distinguish between these photons - the record - and photons from other sources of light (stars etc) that are noise; what makes the photons of the CMBR informative is that they can be reliably linked to the big bang and cannot be explained by the other sources of light. By identifying this, and making the distinction between these photons and photons from other sources we effectively describe what the record is. The CMBR is not a record that Albert's account can easily deal with. It is not clear what the ready state would be in this case - the photons did not exist prior to the big bang and were not knocked out of their expected evolution by it. The definition in terms of noise, however, captures how we distinguish between informative features of the world and background noise. The former are defined in terms of their robustness to the latter.

This is an example of the sort of records we might find in a scientific context, as a more everyday example we can consider a footprint in mud. We distinguish between the outline of the foot and the general unevenness and churned up surface of the mud. The latter acts as noise; when the outline is sufficiently robust and stands out against this background then we have an informative record of a person walking through.

Large interactions that significantly affect the state of the record do not count as noise and if these occur then the original correlation to some past fact can be disrupted, replaced, or destroyed.[7] 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 (the CMBR is a record of the big bang but also acts as noise in telescopic imaging). One might prefer to characterise robustness in terms of stability under perturbations so as to focus on the interactions themselves, and see noise in a signal as the result of these perturbations. However, I have chosen to keep the definition in terms of noise as this is in keeping with the signalling literature where robustness against noise is frequently mentioned, and thinking about signalling helps to highlight that records can be seen as systems that encode information in a robust way.[8] What is important for a record is that one correlation is

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

[8]Additionally, noise is a familiar concept in statistical mechanics, where thermal noise is a familiar concept (see section 4.2) and ideas about noise also enter into accounts of statistical mechanical probab-

singled out and the system becomes an informative signal. The interaction creating this correlation is not fundamentally different from other perturbations of the system, and a record is more than just the general stability of a system, although this is an important factor that the following section will explore. Indeed, a record is not even necessarily a single system, as in the case of the CMBR it can be a collection of entities. In this case thinking about perturbations is less natural. We should focus on the correlation with the past and not the system itself.

A record, then, is a pattern in the world that has a robust correlation to some aspect of the past. This pattern, and the robustness of the correlation, is defined in terms of the separation between the informative record and background noise.

3.2 Stability and Noise

I will now look at how the stability of the record system contributes to robustness and how this can be characterised in terms of noise. This strengthens the claim that noise is definitional of records and not just a practicality. Stability allows us to do away with ready states in many scenarios and is a key part of understanding how a record is correlated to the past, making it a more essential part of what a record is.[9] Looking at how records use stability also tells us more about how records relate to retrodiction; a strict distinction between these two ideas - which I argue is overstated - has been assumed and a more nuanced relationship needs exploring.

The ready state method assumes that a specific past state is needed, and that a separate record of this state is required (this is the route to asymmetry in Albert's account). But it is rare that we actually have such an explicit set up when we make inferences. Instead, we normally rely on the stability that many systems exhibit; for many systems we already have extensive knowledge of their typical evolutions. Using this, we are already able to make fairly accurate inferences about the expected past evolution and from this infer the expected present state (absent interactions), hence doing away with the need for a ready state.[10] 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

ilities (Strevens 2021). This makes the connections to statistical mechanics easier to see.

[9] This is not to say that Albert's account is not helpful in any circumstances. Merely that ready states are not always necessary and often not used at all. They remain a useful tool to understand measuring devices and certain types of records that closely mirror that process.

[10] Again, see fn. 5, we can use a theory of inference to look more closely at what is going on here. My focus on this paper is primarily on distinguishing records from noise, and the theoretical importance of this, rather than on the details of the inferences involved.

specific time in the past where the branch was unbroken but relied on the general stability of the system and our expectations of it.[11]

The role that stability plays in inferences feeds into the distinction between records and retrodiction. Both Albert (2000) (and to a lesser degree Wallace (2013)) claim there is a sharp difference between these two methods of learning about the past. Records have a direct correlation to a specific past state (i.e. we can make a direct inference so some past fact) and we can largely ignore the evolution of the system, 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.

However, looking at retrodiction in a more practical light reveals that this sharp distinction is misleading. The most precise possible retrodiction is evolving the microstate of the world backwards. But such a strict retrodiction is rarely, if ever, performed. Instead, most instances of finding out about the past using dynamical evolution use higher level dynamics or restrict attention to idealized, isolated subsystems. In a less rigorous setting, we 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 rely on information about the dynamical evolution of the system, and do not involve a specific interaction or past state. 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.

When we consider these more common and realistic examples of retrodiction, using generic patterns of dynamical evolution, the distinction between records and retrodiction blurs.

[11] 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. 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 layperson these would be meaningless.

The method by which records give information is not dramatically different from retrodiction, in fact many would consider methods such as historical inference to be examples of records.[12] Brown leaves on the ground appear to be a record of green leaves on the tree despite this inference being basically no different than a retrodiciton. This involves general background knowledge, but nothing as specific as Albert's ready states. Records that involve reference to a specific interaction, which knocked the expected evolution off course, are a special case of a more general idea of records, which involves simply a *robust* correlation to some more generic feature of the past.[13]

Eroding this distinction has implications for what is needed for an explanation of the record asymmetry. Albert's explanation of thermodynamic asymmetry in general relies on the past hypothesis. However, the fact that we have informative records is a precondition for explaining thermodynamic asymmetry. Albert uses the need to explain how we get records as justification for introducing the past hypothesis in the first place. This comes down to the reversibility problem in statistical mechanics (see fn. 3), which tells us we should apply the same reasoning that leads us to predict a higher entropy future to the past, meaning that our records of a low entropy past are incorrect. We instead choose to believe our records over this approach to retrodiciton and introduce the past hypothesis to make sense of why that retrodiction is wrong. If we acknowledge that Albert's idea of records is just a special case of a much more general phenomena, which is closely tied to retrodiction, then explaining records becomes continuous with, although not identical to, explaining the asymmetry of our dynamical laws (most examples of stable dynamics are higher level, macroscopic laws such as thermodynamics that already exhibit asymmetry).[14] Section 5 will continue this line of thought.

Stable and easily predictable dynamics allow us to make quick and easy inferences; but this does not mean that all examples of stable dynamics necessarily produce records, often extra elements are needed. As already stated, it is vague whether we call many sorts of inferences records or retrodiction. Albert's account of an interaction knocking the expected

[12]The various different accounts of records beyond Albert's account do not seem to agree on what exactly counts as a record or not. Reichenbach's account or fork asymmetry as considered by Stradis (2021) would certainly count historical inferences as types of records.

[13]Note that by robust correlation I do not mean a process that always produces the same effect (i.e. something like a robust link between cause and effect). Robustness is meant as it was spelled out in the previous section, it is that we can single out a clear connection to the past against a background of noise and other possible factors that could have contributed to the state of the record. For example, green leaves on trees do not always lead to brown leaves on the ground, the leaves may be blown away or eaten by caterpillars for instance. But when the brown leaves *are* there we can make the inference that they were once green leaves as this is a stable behaviour of this type of system.

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

evolution of a system off course is a definite example of a record and remains a useful tool for analysing how we make inferences in those circumstances. Stable dynamics replaces the need for a ready state to work out the expected evolution but does not change the rest of the reasoning in these circumstances. But Albert's account does not cover all examples of records, ones that are closer to retrodiction (such as the example of inferring from leaves on the ground to leaves on the tree) rely much more heavily on just stable dynamics alone. Overall, examining the role of stable dynamics in records (a feature that all types of records share) rather than looking at ready states (which only some records seem to use) is a more promising route for understanding records and their asymmetry. The next two sections will explore how defining records in terms of their robustness against noise reveals useful theoretical features.

Finally, before moving onto the philosophical implications, this aspect of records can also be understood in terms of robustness against noise: stability implies that a system is not susceptible to noise and a clear distinction between pattern and noise (as in the previous section) can be drawn. Noise can be the result of external influences such as surrounding electromagnetic fields, light pollution etc; but it can also be internal factors in the system such as thermal noise 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, which are generally irrelevant to the overall evolution of the system. All these factors 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 use of the concept as discussed in section 3.1. These sorts of stable patterns that are describable under their own simplified dynamics are exactly the sorts of patterns he identifies as helping us with everyday predictions.

4 Redundancy

With this new account of records in hand I now turn to some of essential features of records that it brings to light, and the philosophical implications of these features. In this section I will look at one particular method of ensuring robustness against noise: *redundancy*. First, I will look at a practical case of how redundancy is used in computing to increase the reliability and robustness of a signal. Then I will extend this to show how redundancy - and robustness in general - can help us understand what is going on in the debate around Maxwell's demon. This gives an example of how redundancy, and the consideration of noise in records, has theoretical as well as practical importance.

4.1 Using Redundancy: Error Correction Coding

The most common method for increasing robustness against noise, beyond the general stability of a system, 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, but 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. A single bit in the state 0 or 1 is replaced by trios of bits in states 0,0,0 or 1,1,1 (I am using 3 bits for proof of concept, in practise it is more complex and depends on the likelihood of errors and on what sort of errors you need to protect against). For simplicity, we assume an error on just one of the bits, which creates states such as 0,0,1 or 1,0,1. These can be distinguished and are then mapped back to the original states.[15] Quantum error correction is more challenging due to the no-cloning theory that prohibits duplication, but the same principles can be applied to produce an error correction theory that makes use of entangled systems (Shor 1995; Nielson and Chuang 2000).

A single state can never be a reliable source of information as it can be changed by errors; we cannot distinguish between correct and incorrect cases. As a result, even in the instances where no errors occur and the system is correctly correlated, we cannot use it as a record

[15] For more details on error correction coding and the array of methods available see Moon (2020). There are many different ways of introducing redundancy. It is also worth clarifying that the word redundant here should not be taken to mean unnecessary or that we can do without it. Redundancy is necessary for error correction, however it is redundant in the sense that the extra bits are not needed just to encode the information in an idealised setting with no noise. Similarly this section will explore how redundancy is necessary to ensure the robustness of records in non-idealised, noisy situations.

because we cannot reliably identify that this is the case. 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 a robust correlation (i.e. the connection becomes robust enough that there is at least a strong likelihood of our inferences about the past based on it 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.[16] 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.

Redundancy is a good method to understand how to deliberately design small scale record systems. But it can also be applied more generally to understand all sorts of records. 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. On a much larger scale, another example of redundancy in a record is the cosmic microwave background radiation. Any individual photon cannot act as a record and its origin cannot be easily traced; even a small sample of the CMBR may not be very informative. However, the distribution of photons across the observable universe is hugely informative, and encodes information that goes beyond what a single photon could ever provide.

4.2 Redundancy in Theory: Maxwell's Demon and The One Molecule Gas Memory Device

The one molecule gas memory device, from Szilard (1929), and the surrounding literature on Maxwell's demon makes clear how the use of redundancy can have philosophical implications and is not simply an operational concern. It also shows how redundancy fits into a formal discussion of records in statistical mechanics. The points discussed here support the broader conclusion from John Norton, who argues that any thermodynamically reversible computations attempted on systems on molecule scales will inevitably be disrupted by thermal fluctuations, and that Maxwell's demon fails on this account without having to rely on more contentious solutions such as Landauer's principle (despite this

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

being the most common resolution cited, e.g. Robertson (2021)). I will show that this follows straightforwardly from the definition of records themselves, and shouldn't be thought of as just a pragmatic limitation. The presentation here will be only a very brief foray into an extensive and complex literature, I will focus just on the particular example of the recording device that is used in this literature. A full, in-depth presentation of the arguments can be found in Norton (2013) (and references therein) and a history of the debate in Norton (2017b) and Robertson (2021).

Maxwell (1871) was considering whether one could control individual particles to create a violation of the second law of thermodynamics. He designed the following thought experiment: the demon is a creature that can observe the exact state of molecules in a gas contained in a box. The box is split into two parts by a partition with a trap door in it. When the demon sees a fast moving particle it lets it through the trap door into one side of the box, and likewise slow moving ones are let into the other half. By doing this the demon is able to increase the temperature of one half of the box in violation of the laws of thermodynamics. Such a demon seemed outside of the boundaries of physics as no creature could have such knowledge or control over individual particles (although modern technology is challenging that (Robertson 2021)), however the problem becomes more concerning when we consider fluctuations. Fluctuations cemented the possibility of short-lived microscopic violations of the second law and did much of the work that the demon was supposed to do, the question now was whether an intelligent agent (like the demon) could trap and accumulate fluctuations into a macroscopic violation. Many new suggestions were put forwards for how to devise a device that could do this (see Norton 2017b).

Coming out of this arose the literature from Szilard (1929) and Landauer (1961) who proposed that, while it may be possible for an intelligent agent to take advantage of microscopic fluctuations in this way, any decrease in entropy that was produced would be cancelled out by thermodynamic costs arising from the operations of the agent itself; so no overall violation of the second law can be achieved. This is now frequently given as the solution to Maxwell's demon. This particularly focused on the fact that the agent would have to measure the system to determine the fluctuated state, record the information so as to act on it, and erase the information to complete the cycle. Szilard argued the thermodynamic cost comes from measurement, Landauer argued it comes from the erasure stage; both make a connection between thermodynamics and information processing.

The basic procedure from Szilard was a process wherein we have a device that fluctuates to a slightly lower entropy (for example the gas in a box fluctuates to one side, effectively decreasing the volume of the gas and hence its entropy[17]). The demon traps it in this

[17]The gas can be just a single molecule so as to make such fluctuations feasibly likely. Macroscopically

state, and then repeats this process indefinitely many times to get a significant violation of thermodynamics. To do this the demon measures the state the gas has fluctuated to, records it, acts, and then erases the information to complete the cycle and start again. To make this analysis possible the demon is idealised as a simple computational device so that thermodynamic analysis can be done on it (Ladyman et al. 2007). The most important condition constraining this is that it must operate in a thermodynamically reversible way; this limits any thermodynamic costs to the logical operations under scrutiny alone (measurement and erasure). Extensive attempts to prove either Szilard's or Landauer's principle have been made, based around this simple process (Ladyman et al. 2007; 2008; Ladyman & Robertson 2014 among others).

The part of this process that is of interest to us here is the memory device that the demon uses, which contains a record. The device standardly used for this is the one molecule gas memory device: we have a single molecule in a box, a partition is then inserted trapping it in one side or the other. Which side the molecule is trapped in encodes a 0 or 1 value respectively, meaning it can encode a single bit of information - a record of what state the system has fluctuated to.[18] The thermodynamic costs of manipulating the device (inserting or removing the partition, or detecting which side the molecule is on) can then be examined (see Norton 2017b; Ladyman & Robertson 2014).

However, Norton (2011; 2013; 2017a, b) argues that this analysis overlooks the thermal fluctuations of the device itself that make any such manipulations impossible (and therefore invalidate these proofs). The main point of the argument is that for these operations to be thermodynamically reversible the device must remain in a delicate equilibrium throughout with no residual energy dissipation. This means that the device will be extremely sensitive and will have thermal energies comparable to the thermal energy of the molecule itself. As such, it too will fluctuate within the exact energy range being studied, which in turn makes smooth, reversible operations impossible. For example, the partition separating the box will fluctuate in position and will not be able to accurately separate the device into two memory states.

The most recent development has been to develop a proof that does not neglect these fluctuations. Myrvold (2020) builds on the earlier proofs mentioned above to do this, he also shows how the proof can be adapted for a quantum mechanical treatment. To deal with fluctuations he uses expectation values, calculated from the probability distribution over the possible states the system fluctuates between, rather than specific values. Any

describing a single molecule as a gas allows us to assign it properties such as entropy and volume.

[18]This strongly resembles Maxwell's original scenario with the gas in a divided chamber. This model is also sometimes called Szilard's one-molecule engine and can itself act as the system in which the thermodynamic violation is induced as well as being the memory device of the agent. See fn. 17.

fluctuations are balanced out by taking the average expected value. This is enough to derive Landauer's principle as a statistical limit. But this analysis has limitations and we can question what exactly has been proved. Although it has established an undeniable link between thermodynamics and logical operations in computing, it is not clear that it has established anything about memory devices in particular, nor does it establish that Landauer's principle is what prevents Maxwell's demon. Myrvold and Norton (2023) have recently clarified that Myrvold's analysis only accounts for fluctuations in the system (including the memory device) but not the apparatus used to control it. Overall, once we include the apparatus, the process is not thermodynamically reversible. So the response to Maxwell's demon that comes out of this is that is it not solely Landauer's principle that thwarts the demon but that the thought experiment fails because it is impossible to have a thermodynamically reversible demon.

I now analyse this in terms of the account of records given here. Norton's arguments focused on the practical difficulties of performing these operations while those attempting to prove Landauer's principle largely put aside practical concerns to focus on conceptual limits, with the expectation that practical considerations are solvable in principle even if we lack the current technology to do so (Norton (2011) highlights this tension). Myrvold attempts to find a balance between the two. But we can go further in arguing that the practical limits become conceptual ones, and that the demon is impossible because it is *conceptually* impossible for it to operate in a thermodynamically reversible way. Under the account given here, it is definitional of a record that it is robust against noise. The one-molecule gas has no such robustness. Thermal noise means that there is no way to reliably create a robust one molecule system with a clear correlation to what we are trying to record. If we attempt to put the device in either the 0 or 1 state there is no way to tell whether we have been successful (that the record is accurately correlated to the desired information) or whether thermal fluctuations have spuriously altered the state. The one-molecule gas device is not a record at all. This is not just a concern over the use of the device, but a limit on whether the required robust correlation can form. To make it into a record we have to use operations that have the required stability and level of control necessary - these are not thermodynamically reversible. The demon requires a record, yet it cannot create one using thermodynamically reversible processes alone.

We can also identify the use of redundancy to create robustness in the way that Myrvold responds to the issue of fluctuations. Redundancy can be identified in the use of expectation values and probability distributions. These can be taken to represent the statistical result of a long sequence of trials or a set of identical memory devices all undergoing the same procedures. Over this sequence the systems will fluctuate into the full range of possible states. Statistically the majority will be affected by minor fluctuations (which are more likely than major ones) that do not alter the overall results being recorded. These states will be correctly correlated to the information we are trying to record. Those displaying a different result to the rest - because they are affected by a major fluctuation - are discounted as errors in an analogue to the error correction coding discussed in the previous section. Taking the expectation value of the probability distribution selects the correctly correlated result. None of the individual systems are informative - any one of them could have been affected by noise, we cannot tell on an individual basis - but the ensemble as a whole is robust against noise and therefore can act as a reliable record.[19]

It is only by characterising records in terms of their robustness against noise, rather than the minimal condition of Albert's ready states, that we see the full complexity of what goes into creating a record. This is a much more stringent requirement and forces us to be far more careful in how we model agents such as Maxwell's demon. 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; similarly it is not that the operations of the demon are constrained by practical considerations but that it is impossible to have thermodynamically reversible demon operating at all. The account of records given here diagnoses the evolution of the debate around Maxwell's demon and Landauer's principle, and how records are used throughout this literature. It shows that the robustness of records against noise is an essential part of understanding how records operate and accounting for it is a central philosophical problem. It also gives a response to the demon that does not depend on this contentious principle but is founded on what it means to have a demon equipped with a memory device.

5 Records, Coarse-graining, and Asymmetry

The main question asked about records is why we have records of the past but not the future; where does this asymmetry come from? The literature on Maxwell's demon is centred around the asymmetry of thermodynamics, but it does not look specifically at the asymmetry of records, for this we must turn to the wider literature on statistical mechanics. Albert's solution to the asymmetry of records was the chain of ready states that links back to the initial state of the universe. This reliance on the past hypothesis fits records into the story he tells about the emergence of asymmetric thermodynamics from

[19]Myrvold includes a note that he does not want to commit to a frequentist interpretation of probabilities, which is what is implied here by transitioning between probability distributions and a redundancy of identical systems. He argues that expectation values can be applied to a single system to derive statistical limits on thermodynamic operations, which do not hold absolutely for any given system but are a useful guide. I argue that, while this is useful for proving theoretical limits, when it comes to actually creating a record in the world this redundancy must become literal and multiple systems must be used to ensure robustness.

statistical mechanics, which is governed by reversible microdynamics, and connects it to all the other temporal asymmetries he argues reduce to statistical mechanics.

The account of records given here, however, offers a different story of how records are connected to the past hypothesis, and temporal asymmetry in general. Robustness against noise implies that modelling records is best done by using methods such as coarse-graining, statistical sets and other ways of defining macroscopic properties. These methods are the basis for Wallace's (see Wallace (2013) for details and a comparison to Albert) alternative account of explaining macroscopic asymmetry within physics. Wallace aims to justify why we have such successful macrodynamics throughout physics that can reliably be used to make predictions about the future, but cannot be reversed and applied to the past (statistical mechanics is the most well-known example of this). His argument runs as follows: for certain systems we can coarse-grain the microstate and evolve the resulting macrostate under macroscopic dynamics. This gives the same final state as evolving the initial microstate under microdynamical laws and then coarse-graining at the end. This implies that the coarse-graining does not remove information that is relevant to the dynamical evolution of the system. The reason that we can do this for dynamical evolutions forwards in time but not in the reverse comes down to the past hypothesis, which for Wallace is the condition that the initial state of the universe was simple and did not contain the sort of microscopic correlations that could become dynamically relevant; this is a very different usage of the past hypothesis compared to Albert's account and Wallace uses it to ground a wide range of statistical methods in physics.[20] Showing how records rely on the same techniques allows us to directly apply much of this argument to records. This connects records to the past hypothesis, as Albert's explanation of asymmetry does, but it a more general way. And it does not rely on the infinite regress of ready states that plays no role in our actual use of records (as discussed in section 3.2).

Leading on from the previous section, redundancy is one feature of records that can lead to the use of statistical methods. Redundancy comes from the fact that the state of any individual system is not robust but a collection of them can be used to identify and correct errors due to noise. Taking the dominant result from a set and disregarding erroneous systems within it is one way of smoothing over microscopic variation. As discussed in section 4.2, Myrvold's use of expectation values and probability distributions are a concrete example of this. These are classic examples of statistical methods. The microscopic details that earlier proofs of Landauer's principle relied on are too susceptible to noise to be robust and these statistical methods are necessary to accurately model the record system.

[20]In comparison, Albert's version of the past hypothesis is a condition that the initial state was low entropy, combined with the statistical postulate which gives a probability distribution over the possible initial microstates of the universe. Wallace (2013) lays out the difference between his past hypothesis and Albert's.

This idea can be generalised. Any instance of using records is going to have to account for the need of robustness. The precise microscopic state is going to be unavoidably susceptible to noise, so to be robust the record must be abstracted from this level of detail.[21] As such, we will use some sort of coarse-graining to capture this - although exactly how we do this will depend on the sort of system.

The definition of noise given in section 3.1 made use of Dennett's (1991) way of distinguishing a pattern from background noise. Records are robust patterns that come with a clear connection to some aspect of the past. Wallace also makes extensive use of Dennett's ideas in how he thinks about the higher level ontologies that we can define based on macroscopic dynamics. Writing down a coarse-grained description of a system is a formal way of distinguishing between a pattern and background noise and connecting it to a simplified macrodynamics that is more predictively useful than the microdynamics. Of course, just distinguishing between pattern and noise is not enough to explain asymmetry. Records are patterns that have a specific correlation with some aspect of the past; the asymmetry of this correlation is what needs explanation. Similarly for Wallace, coarse-graining is only a tool used to develop macrodynamics, the success of those macrodynamics for prediction (but not for retrodiciton) is then based on the past hypothesis. The past hypothesis explains why it is that the details thrown away by coarse-graining are not relevant for the dynamics of the system going forwards in time (but are relevant going backwards) (Wallace provides a far more in-depth account of this that I will not attempt to cover in full detail here). It is harder to apply this to records. Records do not have clearly defined dynamics; they can occur in all kinds of different systems with different physical laws describing them. As such, we cannot give a single account of how records come to have a distinct correlation to some aspect of the past and what dynamics they rely on. The way that records incorporate stability, as discussed in section 3.2, goes some way towards this. Stable dynamics contribute to the robustness of the record and are also the basis on which we make the inference between the record and whatever past fact it is linked to. The sort of stable dynamics that records rely on are in many cases exactly the sort of dynamics that Wallace's account is aimed at explaining. So records may simply inherit their asymmetry from the dynamical systems that exhibit the sort of robustness needed.[22]

Exactly how records inherit their asymmetry from the dynamics they rely on still needs

[21] Although, noise can also include macroscopic perturbations as well. Macroscopic perturbations are more likely to destroy the records (and therefore would not count as noise), but not necessarily. Noise affecting the microscopic state of the record is going to be the easiest to achieve robustness against.

[22] An additional complication to consider is that many records systems may rely on higher level laws from sciences such as biology. These cannot necessarily be put easily into the dynamical forms that Wallace's account analyses. It may be that looking at the asymmetry of causal correlations is of more use here and an account of how these link to Wallace's programme would need to be spelt out.

to be examined. However, identifying the use of methods such as coarse-graining to model records is the first step to doing this and provides a clear link between records and Wallace's account of asymmetry. A more in-depth story of the asymmetry of records must be told on a case by case basis where we can examine specific examples of dynamical systems and look at how distinct correlations between records and the past are formed. The one molecule gas memory device is one example. Another potential case that future work could explore is the use of records in quantum decoherence. Coarse-graining in decoherence, why it is successful and how it leads to asymmetry, has been extensively explored (and forms a large part of Wallace's original arguments). Certain accounts of decoherence have made use records; Zurek (2009)'s quantum darwinism uses the creation of *redundant records* in the environment as a mechanism for decoherence that selects a pointer basis and allows multiple observers to compare measurement results. This is also related to the decoherent histories picture and the coarse-graining used there to derive classical dynamics (Gell-Mann & Hartle 1993). The specific use of redundancy makes immediately clear the similarities with the account of records presented here. A full consideration is beyond the scope of this paper, but decoherence could be a fruitful place to further explore how records use coarse-graining and how this relates to their asymmetry.

Considering coarse-graining and other statistical methods used to model records 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 the reliance of records on redundancy and stability leads to the use of these methods serves as a starting point to understand why and how records are being used throughout physics, and why we have a distinctive record asymmetry. Defining records in terms of their robustness against noise, rather than dismissing this as a merely operational concern, is essential if we are to recognise foundational importance of these features of records.

6 Conclusion

In summary, records can be defined as robust systems that have a clear and identifiable correlation to some aspect of the past. 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. To achieve robustness, records have features such as redundancy, and rely on the stability of the background dynamics. This description of records explains how they are used day to day and takes into account the practical limitations they face. It also identifies two key features: redundancy and the use of statistical methods such as coarse-graining. These previously overlooked features are pivotal to understanding the record asymmetry. They are not operational concerns but are an unavoidable part of how we model records theoretically. They allow us to explain first, why Maxwell's demon fails, and second, how records are connected to Wallace's programme for explaining how the past hypothesis grounds the success (and the asymmetry) of the techniques used in the physics of macroscopic phenomena.

There is still significant work to be done to understand exactly why we have such abundant records of the past but not the future. But focusing on the role of coarse-graining, as a way of capturing the robustness of records against noise, opens up many lines of enquiry and connects records to the significant literature already being developed on understanding these techniques.

References

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

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

Earman, J. (1974). An Attempt to add a little direction to "the problem of the direction of time". *Philosophy of Science,* 41(1), 15-47.

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

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. (2023). Reflections on the asymmetry of causation. *Interface Focus*, 13(3), 20220081.

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.

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

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.

Maxwell, J.C. (1871) *Theory of Heat*. London: Longmans, Green and Co

Moon, T. K. (2020). *Error correction coding: mathematical methods and algorithms*. John Wiley & Sons.

Myrvold, W. C. (2021). Shakin' all over: Proving Landauer's principle without neglect of fluctuations. Forthcoming in BJPS; online at https://www.journals.uchicago.edu/ doi/10.1086/716211.

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

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

Norton, J. D. (2021). *The material theory of induction* (p. 680). University of Calgary Press

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

Robertson, K. (2021). The demons haunting thermodynamics. *Physics Today*, *74*(11), 44-50.

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

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

Strevens, M. (2021). Dynamic probability and the problem of initial conditions. *Synthese*, *199*(5), 14617-14639.

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

Uffink, J. (2002), Essay review: Time and Chance, *Studies in History and Philosophy of Modern Physics* 33, 555-563

Wallace, D. (2013). *The emergent multiverse: Quantum theory according to the Everett interpretation*. Oxford University Press, USA.

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. (2024). Real Patterns in Physics and Beyond.

Wang, Y. (2012). Quantum Computation and Quantum Information. *Statistical Science*, 27(3), 373-394.

Zurek, W. H. (2009). Quantum darwinism. *Nature physics*, 5(3), 181-188.