A Non-Idealised Account of Records

Lucy Mason

Abstract

Records of the past are a pervasive feature of our world which are referred to in many areas of both physics and philosophy, but there is no widespread consensus on what records are. 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 that they must be treated macroscopically. These features have implications the use of records in quantum mechanics and thermodynamics and shows how the misuse of records has caused problems in the philosophy of physics literature.

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 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 which, 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. As such a generalised account of records is needed to understand what a record is and how they play a role in our theories. 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 which 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 centring the requirement that records must be *robustly correlated* with the past state of some other system. 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 dependency between the current state of the record and a specific past event. This correlation is robust; where robustness is against the influence of noise. 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 attention given to records so far has mainly followed Albert in considering primarily their connection to time asymmetry. 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, as does Ismael (2023). I will mainly set aside this time-related literature, although the account proposed here has implications for it which I will discuss, and focus only on one aspect coming out of it: the commonly made assumption 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.[1] 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. It also suggests additionally avenues for exploring the record asymmetry such as a connection to the irreversibility of decoherence and coarse-graining methods used to derive asymmetric macro-dynamics.

However, the use of records goes far beyond questions about asymmetry. Memory systems are used in a wide variety of ways. 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. Additionally records feature in both quantum darwinism and the consistent histories account of the emergence of classicality from quantum mechanics. In these areas that a non-idealised account has great explanatory potential relating to the use of redundancy and the fact that records sit on the boundary between microscopic and macroscopic. I will briefly discuss all

---

[1] 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.

of these, in particular looking 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.

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.[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.

Records encode information about a past event, this is 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 and 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. 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. 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

---

[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. 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 which has many similarities to Albert's.

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.[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 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 many measurement devices designed for the purpose try to achieve as high a level of isolation as possible – there is no way to completely prevent interactions with the many, overlapping systems which 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 which 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 which 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 which 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. 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] In the context of Albert's work this is understandable. His focus is on proving records reliable in the face of the *reversibility problem* and not on their everyday use. 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 but not a direct part of our reasoning.

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 effected 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, 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 3.1 gives a more precise definition of noise and looks at how robustness makes a correlation identifiable. Section 3.2 and 3.3. looks at case studies to show how redundancy is used to protect against it and how this is has theoretical, not just practical, implications.

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 which 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 which 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. 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 which 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 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 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.[4] 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. This extended use will be returned to in section 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 which exist objectively, regardless of our ability as people to read and interpret them. We can only read them *because* there are preexisting structures in the world that make it possible.

Large interactions which 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 which the system takes part in.[5] 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

---

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

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

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 cannot be distinguished. 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 two case studies which show first, how redundancy is used, and second, how this is important for our theoretical use of records.

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 this 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 which 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.[6] 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 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.

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

---

[6] 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.

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. This has lead to the 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) which make reversible thermodynamic processes at this scale impossible. 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 the lack of proper attention to what it takes to be a record led to unnecessary confusion.

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 (2020) 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 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. This isolates the logical operation that takes in place in the system which 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 correctly 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 correlation are exactly those which 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. The fluctuations inherent to any actual processes means a record cannot be created in a thermodynamically reversible way. In terms of proving Landauer's principle as a minimal cost to logical operations this is not too important and 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 abilities of an intelligent agent this is significant as the idea that this is a reliable memory device is central. 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.

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.

4.1 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 which 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.[7] 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. By using this we are already able to make fairly accurate inferences about the past using methods. 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 which 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.[8]

---

[7] 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.

[8] 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

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 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 which make this highly impractical. Instead we tend to use an idealized, isolated subsystem which 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" which he calls historical inference.[9]

---

meaningless.

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

Initially, records are clearly distinguished from the process of retrodiction. 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. Techniques such as historical inference and methods using causal dependencies sit somewhere along this path.[10] And all of these methods 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 which 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 retrodiction is considered above 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 for 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.[11]

---

[10] The various accounts of records 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 likely count historical inferences as types of records. It has no significance where we start applying the label records.

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

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 which 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 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 for these 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. Stability

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 which 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 which are required for records. This use of noise relates strongly back to Dennett's original use of the concept as discussed in section 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.

5 Robust Records must be treated Macroscopically

So what use is this account of records? Section 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 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, however 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 criteria - 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.[12] The need for redundancy in records also increases the size of the system needed 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.[13] 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. 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 some implications here.

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 particularly. Both use records as part of explaining the emergence of the classical world and the irreversibility of decoherence. In quantum darwinism Zurek (2009) describes how the the environment acts as a witness to the state and redundant records of it are made in different parts of the environment. He says that "the proliferation of records enables information about [the system] to be extracted from many fragments of [the environment]...Thus, [the environment] acquires redundant records of [the system]" (Zurek, 2009, pp. 182). The redundancy of the correlations between environment and system is the key to showing how the pointer basis is singled out and how one state can be reliably accessed and agreed upon by multiple observers. The specific use of

---

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

[13] See footnote 1.

redundancy makes immediately clear the similarity between quantum darwinism and the account of records presented here. A full consideration of this is beyond the scope of this paper, but the striking similarities suggests how this account could be used to further our understanding of how records contribute to emergence in this context.

For 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). Understanding records in terms of coarse-graining (and connecting this to the irreversibility of decoherence) 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 which specify an 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.[14] 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 full 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 which records do not obviously have. Case studies such as records in decoherence 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 which 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

---

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

stability requires 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.

6 Conclusion

In summary, records can be defined as systems which enter into robust correlations with the past state of some other system. The robustness criterion is the key feature which leads to the informativeness of the record and allows us to identify the information encoded in the system. The second conclusion of this paper is that robust correlations requires 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 an 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 I hope that 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. However 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.

References

Albert, D. Z. (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.

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.

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

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.

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

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

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.

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.

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

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

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.

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

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.

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

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

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

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. (2012). *The emergent multiverse: Quantum theory according to the Everett interpretation.* Oxford University Press, USA.

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

Wolpert, D. H. (1992). Memory Systems, Computation, and the Second Law of Thermodynamics. *International Journal of Theoretical Physics*, 31, 743-785.

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