

Modeling the Arrows of Time with Causal Multibaker Maps

Aram Eftekar ^{1,*}  and Marcus Hutter ^{2,3} ¹ Independent Researcher, Vancouver, BC V5Y 3J6, Canada² Google DeepMind, London N1C 4AG, UK³ School of Computing, Australian National University, Canberra, ACT 2601, Australia

* Correspondence: aramebtech@gmail.com

Abstract: Why do we remember the past, and plan the future? We introduce a toy model in which to investigate emergent time asymmetries: the causal multibaker maps. These are reversible discrete-time dynamical systems with configurable causal interactions. Imposing a suitable initial condition or “Past Hypothesis”, and then coarse-graining, yields a Pearlean locally causal structure. While it is more common to speculate that the other arrows of time arise from the thermodynamic arrow, our model instead takes the causal arrow as fundamental. From it, we obtain the thermodynamic and epistemic arrows of time. The epistemic arrow concerns records, which we define to be systems that encode the state of another system at another time, regardless of the latter system’s dynamics. Such records exist of the past, but not of the future. We close with informal discussions of the evolutionary and agential arrows of time, and their relevance to decision theory.

Keywords: local causality; Markov property; memory systems; psychological arrow of time; records; second law of thermodynamics; symbolic dynamics

1. Introduction

The “arrows of time” are natural phenomena that are asymmetric under time reversal. They are responsible for the qualitative differences between how we relate to the past and the future, governing our experience of the irreversible passage of time.

To review some important examples, the **thermodynamic arrow of time** is the observation that entropy increases over time [1,2].

The **causal arrow of time** is the observation that causes precede effects [3,4]. It is a key assumption in Bell’s inequality [5,6], and in the design and interpretation of scientific experiments more generally [7–10].

The **epistemic arrow of time** is the observation that there are records of past events, but not of future events [3,4,11,12]. While some authors refer to it as a “psychological arrow” [13,14], the epistemic arrow is not restricted to biological memories: footprints, fossils, impact craters, and photographs all serve as records of past events.

We find other irreversible phenomena in complex, living systems. The **evolutionary arrow of time** is the observation that Darwinian natural selection makes each successive generation better-adapted to its environment [15,16]. At first glance, the evolution of “order” (in the sense of adaptation) seems contrary to the thermodynamic trend toward “disorder” (in the sense of entropy). While the release of entropy into the environment by metabolic processes ensures that there is no contradiction [17], one can still ask why the arrows are aligned.

Finally, the **agential/volitional arrow of time** refers to the observation that intelligent agents model their past as “fixed”, whereas their future is “open” to influence from their present choice of action [3,4,18]. We plan our future retirement, not our past childhood.

The arrows of time are so intertwined, that one may wonder if they are really one and the same. For example, the genes we obtain from Darwinian evolution encode traits that were adaptive in our ancestors’ environment, serving as an epistemic record of past



Citation: Eftekar, A.; Hutter, M. Modeling the Arrows of Time with Causal Multibaker Maps. *Entropy* **2024**, *26*, 776. <https://doi.org/10.3390/e26090776>

Academic Editors: Olimpia Lombardi and Cristian López

Received: 27 May 2024

Revised: 22 July 2024

Accepted: 7 September 2024

Published: 10 September 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

generations. Moreover, evolution produces intelligent agents, who choose a “best” action according to a form of causal reasoning that compares the counterfactual outcomes of different actions [9,19,20]. The past’s immunity to counterfactual change is reminiscent of the robustness of epistemic records, while the future’s epistemic uncertainty is reminiscent of entropy.

Unfortunately, there is not yet any rigorous and general derivation of the arrows of time from fundamental physics. In physical theories with an *initial value formulation*, the Universe is fully determined by its *initial condition* and *dynamical laws* [21]. In this article, we assume that such a formulation exists. Causal reasoning then appears to pose a paradox: interventions and counterfactuals, in describing alternative futures, would violate both determinism and time-reversal symmetry [20,22].

Indeed, the dynamical laws of physics are widely believed to be time-reversal symmetric (more precisely, CPT symmetric [23]). The only remaining piece of our physics formulation, which might produce the needed asymmetry, is its initial condition. It is important to bear in mind that we can equally well determine the Universe by evolving its dynamics backward, from a *final condition* at any later time (i.e., Cauchy surface [21]). Therefore, in order to truly break the symmetry, the initial condition must somehow be special in comparison to later states of the Universe.

The **Past Hypothesis** is the idea that special conditions at the Big Bang explain the arrows of time that we experience today [2,24,25]. A rigorous argument from the Past Hypothesis remains out of reach [12], at least for real systems evolving according to classical or quantum Hamiltonians.

On the other hand, recent advances in Markovian **stochastic thermodynamics** provide a powerful framework for studying irreversible phenomena [26–29]. The idea is to model classical phase space trajectories at a fixed level of precision. The result is a stochastic process over a coarse-grained set of states, each representing some small region in phase space (as in [1] Section 12, or [28] Section 2.7).

We assume this process to be Markovian. While the property of being a Markov chain is invariant under time-reversal, many properties of the Markov kernel are not, such as: homogeneity, locality, algorithmic complexity, and stationarity with respect to a coarse-grained Liouville measure [30–33]. For instance, consider the time-reversed trajectory of a shattered vase: distant shards begin to rise together at the same time, seemingly on their own. At the coarse-grained level, where molecular collisions are hidden, this phenomenon appears nonlocal.

Within the Markovian framework, the thermodynamic arrow is easily shown [27–29]. There is a large literature studying conditions under which the Markov property holds approximately [34–38]; however, these are fairly complicated, obfuscating their connection to the non-thermodynamic arrows of time. Instead, we examine coarse-grainings of simpler toy systems in discrete time, for which the Markov property holds *exactly*.

Gaspard [39] introduces such a system, the *multibaker map*, to study how deterministic and reversible microscopic dynamics can, upon taking a coarse-grained view, emulate diffusive random walks, leading to an irreversible increase in entropy. Altaner and Vollmer [40] introduce an extension, the *network multibaker maps*, to emulate general Markov chains. These maps serve as tractable examples of reversible dynamical systems, on which the theorems of stochastic thermodynamics apply. As such, we can study the relationship between their fine-grained and coarse-grained properties, in the hopes that some insights will carry over to the real Universe.

The Past Hypothesis for a multibaker map is its initial distribution. In [39,40], the fine-grained microscopic details initially are uniformly distributed. An unpublished manuscript [33] relaxes the initialization requirements, by showing that any continuous initial distribution eventually renders the coarse-grained dynamics approximately Markovian, thereby increasing entropy.

While the multibaker maps can model the thermodynamic arrow, the other arrows of time require us to model interactions between multiple physical systems. The reason is

that memory systems need an external system to record, and agents need an environment with which to interact.

Pearl [9] generalizes the Markov property from mere chains to directed acyclic graphs, which model interacting systems. Moreover, he augments the graphs with the *causal* semantics of interventions and counterfactuals. The history of causal concepts spans the disciplines of philosophy, physics, statistics, computer science, and economics, with important contributions from Neyman [41], Reichenbach [3], Lewis [20,42], Bell [5,6], Granger [7,43], Holland [44], Imbens and Rubin [8], and Spirtes et al. [45]. We focus on Pearl [9] due to his convenient structural semantics, which have already been applied to stochastic thermodynamics by Ito and Sagawa [46,47], and Wolpert [48].

In this article, we introduce **causal multibaker maps**. Just as network multibaker maps emulate Markov chains upon coarse-graining, causal multibaker maps emulate Pearlean causal models. In the resulting models, we demonstrate that entropy cannot decrease, and future events cannot be reliably recorded. Thus, we take the causal arrow as fundamental, and use it to explain the thermodynamic, epistemic, evolutionary, and agential arrows. This approach is in line with Markovian stochastic thermodynamics, which uses the Markov assumption to prove thermodynamic relations, but departs from prior efforts to explain the arrows of time, which instead take the thermodynamic arrow as fundamental [3,4,11,12,14,18,49].

Compared to these works, ours has the advantage of being mathematically precise, and applying to a wider variety of coarse-grained dynamics. The drawback is that our underlying microscopic dynamics are of the multibaker type. As such, we can make no definitive conclusions about the arrows of time in the real physical Universe, leaving such extensions to future work.

Nonetheless, to our knowledge, our causal multibaker maps are the first deterministic time-reversible models, that are rigorously shown to possess an emergent causal and epistemic arrow of time at the coarse-grained level. We propose them as a testbed in which to refine hypotheses about the Universe's time asymmetries.

Finally, we point to some related models. *Symbolic dynamics* is a mathematical discipline that studies chaos and irreversibility, using simplified models similar to the multibaker maps [50,51]. Since our models feature subsystems at discrete sites, evolving in discrete time, they can also be considered a kind of *coupled map lattice* [52].

The unpublished manuscript [33] discusses a cellular automaton extension of the multibaker maps. The automaton's dynamical homogeneity in space and time mimics that of field equations in physics. Like field theories, cellular automata propagate influences at the "speed of light", making it difficult to causally insulate subsystems from one another [53]. In order to more easily model periods of non-interaction, our causal multibaker maps allow non-homogeneity.

Article outline. Section 2 sets up some notation, as well as definitions of information-theoretic quantities. Section 3 first develops causal models, and then their microscopic description in terms of causal multibaker maps. In this context, Section 4 discusses each of the arrows of time in turn. Finally, Section 5 discusses future research directions.

2. Preliminaries

We start with some notation. \mathbb{Z} and \mathbb{R} denote the integers and real numbers, respectively, while $\mathbb{Z}_m := \{0, 1, \dots, m - 1\}$ denotes the first m non-negative integers. \emptyset is the empty set, and $\delta_{x, x'}$ is the Kronecker delta that equals 1 if $x = x'$, and 0 otherwise.

When a capital letter such as X refers to a random variable, its lowercase counterparts x, x' refer to the specific values it takes on. This allows us to write $\Pr(x | y)$ as shorthand for the conditional probability expression $\Pr(X = x | Y = y)$.

In classical information theory [54], the **Shannon entropy** of a random variable X is

$$H(X) := \sum_x \Pr(x) \log \frac{1}{\Pr(x)}.$$

The **conditional Shannon entropy** of X , given another random variable Y , is

$$H(X | Y) := H(X, Y) - H(Y) = \sum_{x,y} \Pr(x, y) \log \frac{1}{\Pr(x | y)}.$$

Finally, the **mutual information** between X and Y is

$$I(X : Y) := H(X) + H(Y) - H(X, Y) = H(X) - H(X | Y) = H(Y) - H(Y | X).$$

3. Causal Dynamics

This section consists of four parts, each describing one kind of model. First, we review general causal models on directed acyclic graphs, summarizing some of the core insights from Pearl [9].

Then, we specialize the models to a coarse-grained (mesoscopic) view of causally interacting physical subsystems. We call this view a **persistent causal model**, since each subsystem has an identity that persists over time, its trajectory being represented by a sequence of vertices in the graph. Similar models appear in the stochastic thermodynamics literature [46–48]; since thermodynamic concepts such as heat and work do not enter our discussion, our definitions are a bit simpler. Unlike dynamic Bayesian networks that change over time [55–57], we use a fixed graph to model the states at all times.

To each persistent causal model with doubly stochastic transition probabilities, we associate a **dual model** whose edges point in the opposite temporal direction.

The previous models all have Pearl’s causal arrow of time baked in. In order to connect them to time-symmetric dynamics, we finally present a fine-grained (microscopic) view of persistent causal models: the **causal multibaker maps**. Its dynamics are deterministic and reversible. Therefore, the asymmetry that emerges in the coarse-grained view can be traced back to special properties of the fine-grained initial condition (i.e., a Past Hypothesis).

3.1. Review of Causal Models

Consider a finite set of random variables $\mathbf{V} = \{V_1, V_2, \dots, V_n\}$, which we identify with the n vertices of a directed acyclic graph. Thus, each $V_i \in \mathbf{V}$ is both a random variable and a vertex. Its **parents**, or **direct causes**, are the random variables $PA_i \subset \mathbf{V}$ that have an outgoing edge directly to V_i . Since the sets PA_i together determine all the edges, we can identify the graph with the pair $\mathcal{G} := (\mathbf{V}, \mathbf{PA})$.

We say the random variables \mathbf{V} satisfy the **Markov property**, or are **Markovian**, with respect to the graph \mathcal{G} , if their joint distribution factorizes as

$$\Pr(\mathbf{v}) = \prod_{i=1}^n \Pr(v_i | pa_i). \tag{1}$$

In the special case where $PA_1 = \emptyset$ and $PA_i = \{V_{i-1}\}$ for $i > 1$, \mathbf{V} is called a **Markov chain**.

The Markov property (1) has a useful graphical characterization. For any disjoint sets of variables $X, Y, Z \subset \mathbf{V}$, we say that X and Y are **d-connected**, given Z , if \mathcal{G} has a path $(V_{p(0)}, V_{p(1)}, \dots, V_{p(m)})$ such that:

- $V_{p(0)} \in X$.
- $V_{p(m)} \in Y$.
- For $0 < i < m$, if $V_{p(i)} \notin Z$, then $V_{p(i)}$ is a *chain* or a *fork* on the path, i.e., the edges are oriented as either $V_{p(i-1)} \rightarrow V_{p(i)} \rightarrow V_{p(i+1)}$, $V_{p(i-1)} \leftarrow V_{p(i)} \leftarrow V_{p(i+1)}$, or $V_{p(i-1)} \leftarrow V_{p(i)} \rightarrow V_{p(i+1)}$.
- For $0 < i < m$, if $V_{p(i)} \in Z$, then $V_{p(i)}$ is a *collider* on the path, i.e., the edges are oriented as $V_{p(i-1)} \rightarrow V_{p(i)} \leftarrow V_{p(i+1)}$.

We say that X and Y are **d-separated**, given Z , if they are not d-connected. Verma and Pearl [58] prove the **d-separation criterion**: \mathbf{V} is Markovian with respect to \mathcal{G} iff for all disjoint $X, Y, Z \subset \mathbf{V}$, d-separation of X and Y given Z implies conditional independence of

X and Y given Z . Note that when Z is empty, d -connecting paths are always of the form $V_{p(0)} \leftarrow \dots \leftarrow V_{p(i)} \rightarrow \dots \rightarrow V_{p(m)}$. Therefore, if X and Y are unconditionally dependent, they must share a common ancestor $V_{p(i)}$ in \mathcal{G} , also called a **common cause**.

We should take care not to use the term “cause” loosely. The graph \mathcal{G} is called a **Bayesian network**, if it is only used to convey *probabilistic* information: namely, that \mathbf{V} satisfies its Markov property (1). \mathcal{G} is instead called a **causal network** if, in addition to the Markov property, it also conveys *causal* information about how this distribution would be modified by structural **interventions**. A Bayesian or causal **model** is the combination of a Bayesian or causal network \mathcal{G} , together with all the factors or **mechanisms** $\Pr(v_i | pa_i)$ that complete the right-hand side of (1).

Pearl [59] defines the intervention operator $do(V_i = x)$, which sets $\Pr(v_i | pa_i) := \delta_{v_i, x}$, effectively disconnecting v from its parents. The remaining factors in (1) are left unchanged. In the deterministic case, where \mathbf{V} takes on a fixed value, Bayesian models are trivial: regardless of the edges \mathbf{PA} , there always exists a factorization of the form (1), with $\Pr(v_i | pa_i) := \delta_{v_i, f_i(pa_i)}$ for some functions f_i . In contrast, *causal* models convey nontrivial information even in the deterministic case: the intervention $do(V_i = x)$ alters each descendant V_j of V_i , according to the functions f_k on the paths between them.

With a little more work, Pearl [9] goes on to define the **counterfactual** variables $\mathbf{V}^{do(V_i=x)}$, whose distribution is given by the corresponding intervention $do(V_i = x)$ on \mathbf{V} . It represents the state that our world \mathbf{V} *would have* taken on, *had* the intervention $do(V_i = x)$ been performed. Notice the time-reversal asymmetry: since V_i is cut off from its past (i.e., ancestors in \mathcal{G}) but not its future (i.e., descendants in \mathcal{G}), only the future is subject to change. Thus, causes precede their effects.

The philosophical and physical meaning of interventions and counterfactuals is hotly debated [9,20,25]. They certainly seem important for everyday reasoning: for example, in criminal justice, the accused might be held responsible if their action produced a less favorable outcome for the victim, compared to the counterfactual result of some alternative action. Thus, we entertain an alternative reality in which the accused behaves differently, perhaps even in a manner inconsistent with their psychological nature, which may well be a function of genetics and upbringing. In other words, a counterfactual refers to something that does not occur and may in fact violate the laws of physics [22].

This is sometimes framed as a paradox of “free will” [22], though we prefer to avoid this ambiguous term. Instead, we think of the accused as a being who has evolved a brain, which follows an algorithm to decide an action. One particularly successful kind of algorithm, which Darwinian evolution might plausibly select for, models many available actions as interventions, estimates their outcomes, and then performs the action whose outcome is most desirable to the agent. Thus, although actions are ultimately predetermined, the most successful agents tend to be ones who consider multiple alternatives [60].

Likewise, it is sensible for our systems of ethics and justice to consider counterfactuals, if their purpose is to encourage or deter the actions of citizens. In the scientific disciplines of evolutionary biology, economics, and artificial intelligence, it is common to model agents as optimizing future outcomes toward some objective. Again, this involves a comparison among interventions or counterfactuals [19].

In general, interventions are useful for modeling exogenous influences on an open system. For example, consider the causal network in Figure 1b. If we do not know about the variables in the top row, then we might model the bottom row as an open system, with an unknown exogenous cause acting on the endogenous variable $S_{2,2}$. In this setting, each possible value of $S_{1,1}$ may be modeled by a *do* operator on $S_{2,2}$. The interventions may even be physically realized, if the Universe contains many copies of the system represented by the bottom row, each exposed to a different exogenous context.

3.2. Persistent Causal Models

Now, we use causal networks to model N physical **systems**, numbered $i = 1, \dots, N$, evolving over a sequence of T **events**, numbered $t = 1, \dots, T$. Each system i is associated

with a countable set \mathcal{S}_i of possible coarse-grained states. In general, any subset A of the **Universe** $U := \{1, \dots, N\}$ may be viewed as a composite system, with associated state space $\mathcal{S}_A := \prod_{i \in A} \mathcal{S}_i$. If $A \subset B \subset U$, we say that A is a **subsystem** of B .

Each event t is associated with a disjoint pair of composite systems $e(t), c(t) \subset U$. $e(t)$ represents the systems that evolve during the event t , while all other systems are held fixed. The evolution is influenced by the systems $c(t)$, so that the joint state probability distribution of the systems $e(t)$ at time t is defined as some function of the state of the systems $e(t) \cup c(t)$ at time $t - 1$. $e(t)$ should be non-empty, since otherwise the event does nothing; however, $c(t)$ can be empty.

Formally, a persistent causal model is a Pearlean causal model whose graph \mathcal{G} has some vertices corresponding to each time $t = 0, 1, \dots, T$. For the initial time $t = 0$, \mathcal{G} has $N + 1$ vertices:

$$\begin{aligned} E_0 &\in \mathcal{S}_U, & PA(E_0) &:= \emptyset, \\ \text{For } i \in U : & S_{0,i} \in \mathcal{S}_i, & PA(S_{0,i}) &:= \{E_0\}. \end{aligned}$$

$\Pr(E_0)$ is the initial state distribution of our Universe, while each $S_{0,i} := (E_0)_i$ is deterministically assigned the initial state of system $i \in U$.

For each non-initial time $t = 1, \dots, T$, \mathcal{G} has $|e(t)| + 1$ vertices:

$$\begin{aligned} E_t &\in \mathcal{S}_{e(t)}, & PA(E_t) &:= \{S_{t-1,j} : j \in e(t) \cup c(t)\}, \\ \text{For } i \in e(t) : & S_{t,i} \in \mathcal{S}_i, & PA(S_{t,i}) &:= \{E_t\}. \end{aligned}$$

In order for the expression for $PA(E_t)$ to make sense, we define $S_{t,i} := S_{t-1,i}$ for the non-evolving systems $i \in U \setminus e(t)$; these are not additional vertices, but rather, aliases for vertices from an earlier time. $\Pr(E_t | PA(E_t))$ is the dynamics of event t , while each $S_{t,i} := (E_t)_i$ is deterministically assigned the evolved state of system $i \in e(t)$.

For convenience, we denote the state of a composite system $A \subset U$ at time t by $S_{t,A} := (S_{t,i} : i \in A)$. Our causal network indicates a joint probability distribution that factorizes as (1). Since $E_0 = S_{0,U}$ and $E_t = S_{t,e(t)}$ are effectively contained in the collection of random variables \mathbf{S} , it suffices to expand the latter's distribution:

$$\Pr(\mathbf{s}) := \Pr(s_{0,U}) \prod_{t=1}^T \Pr(s_{t,e(t)} | s_{t-1,e(t)}, s_{t-1,c(t)}). \tag{2}$$

In summary, a persistent causal model is specified by its event structure $(e(t), c(t))_{t=1}^T$ (which determines the graph \mathcal{G}), its **initial condition** $\Pr(s_{0,U})$, and its events' **forward mechanisms** $\Pr(s_{t,e(t)} | s_{t-1,e(t)}, s_{t-1,c(t)})$. The case $N = 1$ corresponds to a Markov chain, whereas having $N > 1$ allows us to model interactions between N systems.

In cases where the systems are initially independent, we can omit the vertex E_0 and directly assign each respective initial distribution $\Pr(S_{0,i})$:

$$\text{For } i \in U : \quad S_{0,i} \in \mathcal{S}_i, \quad PA(S_{0,i}) := \emptyset.$$

Similarly, for events t that evolve only one system i , we can omit the vertex E_t and directly assign the forward mechanism to $\Pr(S_{t,i} | PA(S_{t,i}))$:

$$\text{For the unique } i \in e(t) : \quad S_{t,i} \in \mathcal{S}_i, \quad PA(S_{t,i}) := \{S_{t-1,j} : j \in e(t) \cup c(t)\}.$$

Thus, in the simple case where the initial states are independent and the systems evolve one at a time, \mathcal{G} has $n = N + T$ vertices: one for the initial state of each system $i = 1, \dots, N$, and one for each event $t = 1, \dots, T$. Let's consider one such example.




Figure 1. (a) The graph of a persistent causal model for two systems. Vertices are drawn as circles containing their variable name; edges are drawn as arrows. System 1 (top row) evolves autonomously as a Markov chain. At time $t = 2$, it influences the external memory represented by System 2 (bottom row). (b) The same model, with alias variables (dashed circles) to represent non-evolving systems. These variables are exactly equal to their predecessor along a dashed edge.

Example 1. From the two-system causal network in Figure 1a, we read off $e(1) = e(3) = \{1\}$, $c(1) = c(3) = \emptyset$, $e(2) = \{2\}$, and $c(2) = \{1\}$. Suppose System 1 has an integer-valued state that evolves autonomously on a random walk: at each of the times $t = 1, 3$, it either stays still, increments, or decrements, each with probability $1/3$. Moreover, suppose that at the time $t = 2$, System 1 reversibly writes its state to an external memory, represented by System 2.

Rather than explicitly tabulate all of the conditional probabilities in (2), we can implicitly define each mechanism in terms of a variable assignment:

$$\begin{aligned} S_{0,1} &:= S_{0,2} := 0, \\ S_{1,1} &:= S_{0,1} + \text{Uniform}\{-1, 0, 1\}, \\ S_{2,2} &:= S_{1,1} \oplus S_{0,2}, \\ S_{3,1} &:= S_{1,1} + \text{Uniform}\{-1, 0, 1\}, \end{aligned}$$

where \oplus denotes bitwise exclusive-or (i.e., binary addition without carries, so $5 \oplus 3 = 6$), and Uniform indicates an independent sample with uniform probability of being each element of the enclosed set. In Figure 1b, we add the alias variables $S_{1,2} := S_{0,2}$, $S_{2,1} := S_{1,1}$, and $S_{3,2} := S_{2,2}$, for the non-evolving system at each time step.

3.3. The Dual of a Persistent Causal Model

In order to reverse the causal arrow, we need an additional property: an event t is **doubly stochastic** if, for all $y \in \mathcal{S}_{e(t)}$ and $z \in \mathcal{S}_{c(t)}$,

$$\sum_{x \in \mathcal{S}_{e(t)}} \Pr(S_{t,e(t)} = y \mid S_{t-1,e(t)} = x, S_{t-1,c(t)} = z) = 1. \tag{3}$$

We then define the event’s **dual mechanism**

$$\Pr(\tilde{S}_{t-1,e(t)} = x \mid \tilde{S}_{t,e(t)} = y, \tilde{S}_{t,c(t)} = z) := \Pr(S_{t,e(t)} = y \mid S_{t-1,e(t)} = x, S_{t-1,c(t)} = z), \tag{4}$$

which is a valid conditional probability distribution whenever (3) holds.

Now suppose we have a persistent causal model (2), whose events are all doubly stochastic (3); and we also have a **final condition** $\Pr(\tilde{s}_{T,U})$. Then, the dual mechanisms (4), together with the aliases $\tilde{S}_{t-1,i} := \tilde{S}_{t,i}$ for $i \in U \setminus e(t)$, determine a dual collection of random variables $\tilde{\mathbf{S}}$, whose joint probability distribution is given by

$$\Pr(\tilde{\mathbf{S}}) := \Pr(\tilde{s}_{T,U}) \prod_{t=1}^T \Pr(\tilde{s}_{t-1,e(t)} \mid \tilde{s}_{t,e(t)}, \tilde{s}_{t,c(t)}). \tag{5}$$

The causal network for $\tilde{\mathbf{S}}$ is inferred from Equation (5). Rather than describe it in abstract terms, we demonstrate with an example.

Example 2. The reader may verify that the events in Example 1 are not only doubly stochastic, but also self-dual. Therefore, its dual mechanisms are well-defined and identical to its forward mechanisms. Moreover, the sequence of events is symmetric, with System 1 taking a random step both before and after writing to System 2. If we complete the symmetry by setting the dual model’s final condition to match the forward model’s initial condition, then the two models become fully identical, except that the dual model’s time indices are reversed:

$$\begin{aligned} \tilde{S}_{3,1} &:= \tilde{S}_{3,2} := 0, \\ \tilde{S}_{2,1} &:= \tilde{S}_{3,1} + \text{Uniform}\{-1, 0, 1\}, \\ \tilde{S}_{1,2} &:= \tilde{S}_{2,1} \oplus \tilde{S}_{3,2}, \\ \tilde{S}_{0,1} &:= \tilde{S}_{2,1} + \text{Uniform}\{-1, 0, 1\}. \end{aligned}$$

The causal network for this dual process, shown in Figure 2, is a mirror image of the original network in Figure 1. However, we must remark that the dual process does not undo the forward process. That is, suppose we change the dual model’s final condition to match the forward model’s final distribution:

$$(\tilde{S}_{3,1}, \tilde{S}_{3,2}) := (S_{3,1}, S_{3,2}).$$

Then, the dual process does not undo System 1’s random walk to return to $S_{0,1} := 0$. Instead, it takes two additional steps, making $\tilde{S}_{0,1}$ the result of four independent random steps; its value is randomly distributed over the range $[-4, 4]$.




Figure 2. (a) The dual of Figure 1a; a causal network whose edges point to decreasing times. (b) The dual network with alias variables (dashed circles) included. Note that dashed edges appear in exactly the same positions as in Figure 1b, but they point to different dashed circles.

Intuitively speaking, the dual process corresponds to a coarse-grained approximation of time reversal, in which a loss of fine-grained correlations prevents random processes from being undone. It would be like trying to unstir a cup of coffee by reversing the mixing spoon’s trajectory. To make this idea precise, we need a fine-grained model.

3.4. A Microscopic View: Causal Multibaker Maps

Our fine-grained model, the **causal multibaker map**, is a multi-system extension of the traditional baker and multibaker maps [39,40]. Its dynamics are deterministic and reversible. It can be given either an initial or a final condition; accordingly, a causal multibaker map’s coarse-grained behavior is either that of the persistent causal model (2), or its dual (5), respectively.

Moreover, the correspondence between fine-grained and coarse-grained dynamics is localized to individual events. Thus, an event’s coarse-grained transition probabilities along the causal network’s edges can be predicted from local knowledge of the event’s fine-grained mechanism, and controlled by modifying this mechanism. In contrast, an event’s transition probabilities against the edges depend nonlocally on other mechanisms. This provides a causal arrow of time.

To describe a causal multibaker map, assume N, T , and $(e(t), c(t))_{t=1}^T$ are given as before, and fix an integer $m > 1$. A fine-grained **microstate** for system i is a pair of the form

$$(x, (\dots, r_{-2}, r_{-1}, r_0, r_1, r_2, \dots)), \tag{6}$$

consisting of a coarse-grained state $x \in \mathcal{S}_i$, along with a bi-infinite sequence of microvariables $r_k \in \mathbb{Z}_m$. The full microstate of our model Universe at any given time consists of N such pairs, one for each system.

To get some geometric intuition, we can identify \mathcal{S}_i with \mathbb{Z} , and interpret (6) as a point in the two-dimensional “phase space” $\mathbb{R} \times [0, 1]$, whose base m representation is

$$(x.r_0r_{-1}r_{-2}\dots, 0.r_1r_2r_3\dots). \tag{7}$$

However, the symbolic representation (6) will be more convenient.

An **initial value formulation** determines the entire trajectory of the model Universe from two objects: (1) an initial distribution over the joint microstate of the systems, and (2) the event dynamics, which consist of deterministic state transformations.

The coarse-grained variables are initialized to a given distribution $\Pr(s_{0,U})$, on which we place no restrictions. Meanwhile, the microvariables are initialized to be jointly independent (of each other and the coarse-grained variables), and uniformly distributed, taking each value in \mathbb{Z}_m with equal probability $1/m$.

All that remains is to specify the dynamics at each event t . It consists of two stages. The first stage is a shift transformation on one system $i \in e(t)$; for concreteness, we let i be the least element of $e(t)$. All of system i 's microvariables r_k shift one position to the left.

The second stage is a bijective transformation

$$f_t[s_{t-1,c(t)}] : \mathcal{S}_{e(t)} \times \mathbb{Z}_m \rightarrow \mathcal{S}_{e(t)} \times \mathbb{Z}_m,$$

that depends on the coarse-grained state $s_{t-1,c(t)}$ of the influencing systems $c(t)$. It is applied jointly to the coarse-grained state $s_{t-1,e(t)}$ of the evolving systems, and the newly centered microvariable r_1 of the system $i \in e(t)$:

$$(s_{t,e(t)}, r'_1) := f_t[s_{t-1,c(t)}](s_{t-1,e(t)}, r_1). \tag{8}$$

In the case where $e(t)$ contains only one system i , its full two-stage transformation is summarized as follows:

$$\begin{aligned} & (s_{t-1,i}, (\dots, r_{-2}, r_{-1}, r_0, r_1, r_2, \dots)) \\ & \xrightarrow{\text{shift}} (s_{t-1,i}, (\dots, r_{-1}, r_0, r_1, r_2, r_3, \dots)) \\ & \xrightarrow{\text{transform}} (s_{t,i}, (\dots, r_{-1}, r_0, r'_1, r_2, r_3, \dots)). \end{aligned}$$

With these deterministic dynamics, the only source of randomness is the initial condition. Since the microvariables are initialized independently and uniformly, we think of each $r_k \in \mathbb{Z}_m$ as a fair m -sided die that can be used to emulate a stochastic transition of the coarse-grained variables. As a result, when we ignore all of the r_k , the coarse-grained variables' trajectory is given by the persistent causal model (2), with the forward-time transition probabilities

$$\Pr(s_{t,e(t)} \mid s_{t-1,e(t)}, s_{t-1,c(t)}) = \frac{1}{m} \left| \left\{ r_1 : \exists r'_1, f_t[s_{t-1,c(t)}](s_{t-1,e(t)}, r_1) = (s_{t,e(t)}, r'_1) \right\} \right|. \tag{9}$$

Bijection of $f_t[s_{t-1,c(t)}]$ implies double stochasticity of the forward mechanism (9).

Conversely, every doubly stochastic persistent causal model, whose probabilities are multiples of $1/m$, is emulated by a causal multibaker map with a suitable choice of bijections $f_t[s_{t-1,c(t)}]$. Indeed, given the desired transition probabilities, we need only assign each pair $s_{t-1,e(t)}, s_{t,e(t)} \in \mathcal{S}_{e(t)}$ to each other with multiplicity $m \cdot \Pr(s_{t,e(t)} \mid s_{t-1,e(t)}, s_{t-1,c(t)})$,

for each fixed t and $s_{t-1,c(t)}$. One way to accomplish this is to fix any total order $<$ on $\mathcal{S}_{e(t)}$. Then, for all trajectories \mathbf{s} , events t , and $r \in \mathbb{Z}_{m \cdot \Pr(s_{t,e(t)} | s_{t-1,e(t)}, s_{t-1,c(t)})}$, let

$$f_t[s_{t-1,c(t)}] \left(s_{t-1,e(t)}, r + \sum_{s'_{t,e(t)} < s_{t,e(t)}} m \cdot \Pr(s'_{t,e(t)} | s_{t-1,e(t)}, s_{t-1,c(t)}) \right) \\ := \left(s_{t,e(t)}, r + \sum_{s'_{t-1,e(t)} < s_{t-1,e(t)}} m \cdot \Pr(s_{t,e(t)} | s'_{t-1,e(t)}, s_{t-1,c(t)}) \right).$$

Double stochasticity (3) ensures that each $f_t[s_{t-1,c(t)}]$ is well-defined and bijective. Moreover, the fact that r takes on $m \cdot \Pr(s_{t,e(t)} | s_{t-1,e(t)}, s_{t-1,c(t)})$ values, each realized with probability $1/m$, ensures the required transition probabilities (2). The Birkhoff-von Neumann theorem implies that $f_t[s_{t-1,c(t)}]$ may even be defined in such a way that $r'_1 = r_1$ always holds in (8); however, such a definition would be more complicated [61].

Thus, under mild restrictions on the persistent causal model, we can always produce a multibaker map that emulates it. The restriction that the probabilities have a common denominator m can be removed, using techniques from [40] or [33]; we only used it here to take advantage of the simpler m -ary shift representation. The other restriction is double stochasticity (3). It comes from the microscopic reversibility of the transformation functions $f_t[s_{t-1,c(t)}]$, and is a special case of a stationarity condition that comes from coarse-graining a broader class of measure-preserving transformations [33,40].

The correspondence between fine-grained and coarse-grained views has the additional feature of being *local*: the definition of a fine-grained transformation f_t uniquely determines the coarse-grained transition probabilities (9), regardless of any other events. So for example, to construct a time-homogeneous Markov chain with $N = 1$, it suffices to repeat the same microscopic transformation f_t at every time step.

This works because the transformation (8) is always applied on a “fresh die” r_1 , which is uniform and independent of the coarse-grained trajectory from initialization up to the present. If instead we view the process in reverse, undoing events in reverse order, then the inverse function $f_t^{-1}[s_{t,c(t)}]$ is applied to the possibly correlated pair $(s_{t,e(t)}, r'_1)$. As such, the time-reversed transition probabilities cannot be locally derived from f_t alone; in general, they depend on the earlier transformations and initial conditions. Rather than follow a locally well-defined statistical law as in (9), the reverse dynamics must “conspire” to return to the initial condition.

Example 3. To get a fine-grained model of Example 1, let $m := 3$. For $s_{t,i} \in \mathbb{Z}$ and $r \in \mathbb{Z}_3$, define the microscopic transformations

$$f_1(s_{0,1}, r) := (s_{0,1} + r - 1, r), \\ f_2[s_{1,1}](s_{0,2}, r) := (s_{0,2} \oplus s_{1,1}, r), \\ f_3(s_{1,1}, r) := (s_{1,1} + r - 1, r).$$

The fine-grained trajectory of the two systems is listed in Table 1. Here we see a thermodynamic arrow of time: each step of System 1’s random walk adds an independent random variable (first r_1 , then r_2), increasing the entropy of its coarse-grained state. We also see an epistemic arrow of time: at times $t \geq 2$, System 2 maintains a record of $s_{1,1} = r_1 - 1$.

Starting from the final state at the bottom row of Table 1, we can follow the trajectory in reverse, applying the inverse transformations

$$f_3^{-1}(s_{3,1}, r) := (s_{3,1} + 1 - r, r), \\ f_2^{-1}[s_{2,1}](s_{3,2}, r) := (s_{3,2} \oplus s_{2,1}, r), \\ f_1^{-1}(s_{2,1}, r) := (s_{2,1} + 1 - r, r).$$

In doing so, we would witness two strange “miracles”. First, going from $t = 3$ to $t = 2$, System 1’s random walk happens to converge to the value recorded in System 2, despite the systems not interacting during this time interval. And finally, going from $t = 1$ to $t = 0$, the random walk converges to its starting point, clearing all of its entropy.

The miracles are explained by the correlation between the “final conditions”

$$\begin{aligned} s_{3,1} &= r_1 + r_2 - 2, \\ s_{3,2} &= r_1 - 1, \end{aligned}$$

and the microvariables (r_1, r_2) that control the randomness of the coarse-grained dynamics.

If instead we set proper final conditions, in which the microvariables at $t = 3$ are independent of $(s_{3,1}, s_{3,2})$, then the arrow of time would reverse. To be precise, the coarse-grained behavior would be given by the dual model in Example 2. Tables 2 and 3 list the fine-grained dual trajectories for both sets of coarse-grained final conditions considered in Example 2.

In general, the causal arrow of time is determined not by the t -coordinate, but by the graph’s edges. Tables 1–3 are all arranged such that the arrow of time points downward.

Table 1. The trajectory of the systems modeled in Example 3. Each $r_k \in \mathbb{Z}_3$ is independent and uniformly distributed.

Time t	State of System 1	State of System 2
0	$(0, (\dots, r_0, \dots))$	$(0, (\dots))$
1	$(r_1 - 1, (\dots, r_1, \dots))$	$(0, (\dots))$
2	$(r_1 - 1, (\dots, r_1, \dots))$	$(r_1 - 1, (\dots))$
3	$(r_1 + r_2 - 2, (\dots, r_2, \dots))$	$(r_1 - 1, (\dots))$

Table 2. The dual model’s trajectory, when its final condition is set to the initial condition of Table 1. Although the expressions appear different, the symmetry of the events renders this coarse-grained trajectory’s distribution identical to that of Table 1.

Time t	State of System 1	State of System 2
3	$(0, (\dots, r_0, \dots))$	$(0, (\dots))$
2	$(1 - r_0, (\dots, r_{-1}, \dots))$	$(0, (\dots))$
1	$(1 - r_0, (\dots, r_{-1}, \dots))$	$(1 - r_0, (\dots))$
0	$(2 - r_0 - r_{-1}, (\dots, r_{-2}, \dots))$	$(1 - r_0, (\dots))$

Table 3. The dual model’s trajectory, when the coarse-grained part of its final condition is set to the final distribution of Table 1. The reason it does not simply retrace the trajectory of Table 1, is that the microvariables here are reset to fresh independent values.

Time t	State of System 1	State of System 2
3	$(r_1 + r_2 - 2, (\dots, r_0, \dots))$	$(r_1 - 1, (\dots))$
2	$(r_1 + r_2 - r_0 - 1, (\dots, r_{-1}, \dots))$	$(r_1 - 1, (\dots))$
1	$(r_1 + r_2 - r_0 - 1, (\dots, r_{-1}, \dots))$	$(r_1 - 1 \oplus r_1 + r_2 - r_0 - 1, (\dots))$
0	$(r_1 + r_2 - r_0 - r_{-1}, (\dots, r_{-2}, \dots))$	$(r_1 - 1 \oplus r_1 + r_2 - r_0 - 1, (\dots))$

We conclude that the causal arrow of time always points away from the initialization time, at which the microvariables are independent and uniform. In fact, the initialization need not be so strict: provided that the initial distribution is reasonably “smooth” on its phase space geometry (7), the microvariables r_k become increasingly uniform as $k \rightarrow \pm\infty$. After sufficiently many iterations of the shift transformation (or its inverse), it follows that the coarse-grained dynamics become indistinguishable from the case of uniform microvariable initialization; see [33] for a rigorous proof. Therefore, as depicted in Figure 3, the arrow of time points away from the initialization time toward $\pm\infty$, except that its direction may temporarily be ambiguous near initialization.

One may speculate about whether the real-world Past Hypothesis (at the Big Bang) can be expressed in similar terms. Instead of the shift transformation, a chaotic Hamiltonian evolution would give the phase space distribution an increasingly fine filamentary structure, until it effectively “looks uniform” on a larger region of phase space [2]. Of course, this argument applies equally well in reverse; Boyle et al. [62,63] suggest that a CPT-inverted dual dynamics may prevail before the Big Bang.




Figure 3. A conceptual visualization of the arrow of time. The horizontal axis represents the time coordinate, while the vertical axis represents entropy. If the state distribution of a multibaker map is absolutely continuous at some initialization time (the small cross), then some duration around it (the wavy central region) may behave ambiguously, but eventually the arrow of time points consistently away to $\pm\infty$. The dynamics of the left region are dual to those of the right region.

4. The Arrows of Time

The causal multibaker maps exhibit some time asymmetries that are analogous to those we see in the real Universe. As such, these maps can be said to model the arrows of time in a mathematically precise and tractable way. We now examine these asymmetries.

4.1. The Causal Arrow of Time

Despite their deterministic and reversible fine-grained evolution, we saw that causal multibaker maps with suitable initial conditions exhibit asymmetric coarse-grained behavior. This behavior is exactly described by the persistent causal model (2). It satisfies the graph’s Markov property, so that, for example, every correlated pair of events has a common cause from an earlier time. Moreover, the coarse-grained mechanisms are local functions of the corresponding fine-grained transformations.

To be precise, when the model represents the Universe as a whole, then it can only be interpreted as a Bayesian model, not a causal one. The *do* operator, required for causal semantics, amounts to a modification of the dynamics, which is not possible from a global point of view. On the other hand, we gave plausible arguments for the *do* operator’s practical relevance, to model exogenous influences on subsystems that occur repeatedly in the Universe.

Indeed, Bell [5,6] observed that the validity of the scientific method rests on an experimenter’s ability to independently choose or randomize control variables in a repeated experiment. Reusing the causal network in Figure 1, we may interpret System 1 as an experimenter, who sets a control variable in System 2. The interaction mechanism $\Pr(s_{2,2} \mid s_{1,1}, s_{0,2})$ is determined by local considerations, represented by the fine-grained transformation f_2 . In particular, $S_{2,2}$ can be randomized independently of past states, regardless of any other mechanisms in the model. The reverse is not possible: depending on other mechanisms, the future coarse-grained state of the Universe may in general correlate with $S_{0,2}$ and/or $S_{2,2}$.

As an additional remark, if we want to define interventions and counterfactuals on a causal multibaker map, there is no conflict between Lewis [20]’s closest-world semantics and Pearl [9]’s structural semantics. Lewis’ version of the counterfactual $\mathbf{S}^{do(S_{t,i}=x)}$ is defined as the Universe closest to \mathbf{S} that satisfies $S_{t,i} = x$. If we define distance by the number of altered fine-grained mechanisms, breaking ties by the number of altered state values, then the closest Universe satisfying $S_{t,i} = x$ is one which alters only the mechanism f_i with respect to system i . Thus, we recover Pearl’s semantics.

4.2. The Thermodynamic Arrow of Time

The second law of thermodynamics can be derived as a consequence of the Markov property and double stochasticity. In the special case of Markov chains, the second law is a very well-known mathematical theorem.

Theorem 1 (Second law of thermodynamics for Markov chains). *Consider any Markov chain, given by (2) with $N = 1$. Suppose $0 \leq t \leq u \leq T$, and that all events $v = t + 1, \dots, u$ are doubly stochastic. Then,*

$$H(S_{t,1}) \leq H(S_{u,1}).$$

Proof. See Section 4.4 in Cover and Thomas [54]. \square

Theorem 1 provides the most well-understood arrow of time. It implies that all entropy-increasing processes are irreversible. Perhaps it is a bit ironic that, for the multibaker maps, the irreversible growth of entropy is a consequence of double stochasticity, which in turn is a consequence of f_t 's reversibility.

By conditioning on any exogenous influences, we extend Theorem 1 to non-isolated systems.

Theorem 2 (Second law of thermodynamics for causal models). *Consider any persistent causal model (2), and $0 \leq t \leq u \leq T$. Let $A, B \subset U$ be composite systems with the property that for all $v = t + 1, \dots, u$:*

- *Either $e(v) \subset A$, or $e(v) \subset B$, or $e(v) \subset U \setminus (A \cup B)$.*
- *If $e(v) \subset A$, then $c(v) \subset A \cup B$ and the event v is doubly stochastic.*
- *If $e(v) \subset B$, then $c(v) \subset B$.*

Then,

$$H(S_{t,A} | S_{t,B}) \leq H(S_{u,A} | S_{u,B}). \tag{10}$$

Proof. Removing elements from A that are shared with B does not change the conditional entropy; hence, we lose no generality in assuming that A and B are disjoint. By mathematical induction, it suffices to prove (10) over the span of a single event $u = t + 1$. Depending on the systems $e(u)$ that it acts on, there are three cases to consider.

When $e(u) \subset U \setminus (A \cup B)$, the composite systems A and B are unchanged, so (10) holds with equality.

When $e(u) \subset A$, we have $S_{t,B} = S_{u,B}$. Moreover, $c(v) \subset A \cup B$, so the composite system $A \cup B$ undergoes doubly stochastic transition probabilities with no outside influence. Treating $A \cup B$ as one system, Theorem 1 implies $H(S_{t,A \cup B}) \leq H(S_{u,A \cup B})$. Therefore,

$$\begin{aligned} H(S_{t,A} | S_{t,B}) &= H(S_{t,A \cup B}) - H(S_{t,B}) \\ &\leq H(S_{u,A \cup B}) - H(S_{u,B}) \\ &= H(S_{u,A} | S_{u,B}). \end{aligned}$$

Finally, when $e(u) \subset B$, we have $c(u) \subset B$ and $S_{t,A} = S_{u,A}$. By the d-separation criterion, the latter variable is conditionally independent of $S_{u,B}$, given $S_{t,B}$, so the data processing inequality ([54] Thm 2.8.1) implies $I(S_{t,A} : S_{t,B}) \geq I(S_{u,A} : S_{u,B})$. Therefore,

$$\begin{aligned} H(S_{t,A} | S_{t,B}) &= H(S_{t,A}) - I(S_{t,A} : S_{t,B}) \\ &\leq H(S_{u,A}) - I(S_{u,A} : S_{u,B}) \\ &= H(S_{u,A} | S_{u,B}). \end{aligned} \quad \square$$

Since the conditional entropy in question can be expressed as

$$H(S_{t,A} | S_{t,B}) = H(S_{t,A}) - I(S_{t,A} : S_{t,B}),$$

Theorem 2 allows the composite system A 's entropy to decrease, provided that the mutual information term also decreases. The influence of a "Maxwell's demon" B , with knowledge (i.e., mutual information) about A 's state, can therefore decrease A 's entropy [46,64].

4.3. The Epistemic Arrow of Time

"The epistemic arrow of time is the fact that our knowledge of the past seems to be both of a different kind and more detailed than our knowledge of the future."
- Wolpert and Kipper [12]

Consider the difference between weather forecasts (of the future) and historical weather records (of the past). Weather forecasts are derived from computationally intensive simulations seeded with precise global measurements, and become unreliable beyond about a week. In contrast, precise records can be produced with minimal computational effort or meteorological expertise, and then maintained for millennia.

Wolpert and Kipper [11,12] distinguish memories that are "Type-2" or "computer-type", from those that are "Type-3" or "photograph-type". A computer-type memory is not considered a record. It uses information about the present state and dynamics of another system, to simulate its evolution very precisely, either backward or forward in time. In this manner, computer-type memories can produce predictions about both the past and future.

In order for the simulation to be accurate, computer-type memories require knowledge of the other system's dynamics, including its interactions with additional systems, which may in turn interact with more systems, and so on. Moreover, when the dynamics are chaotic, the simulation must start from extremely fine-grained knowledge of the state, which is simply not available in our coarse-grained view.

In contrast, photograph-type memories do serve as records. Slightly adapting the definition of Mlodinow and Brun [14], we consider records to be systems, whose coarse-grained state at some designated time t_{read} , is a non-constant function of another system's coarse-grained state, at some other time t_{event} , regardless of the latter system's dynamics. In Example 1, System 2 fits this definition with $t_{event} = 1$ and $t_{read} \in \{2, 3\}$. System 2 does not simulate System 1's dynamics; instead, it stores the event state directly. As a result, even if we were to change the initial state and dynamics of System 1 entirely, System 2 would still successfully record the event $S_{1,1}$.

Records can only be of the past, as any attempt to record a future event would be foiled by some choice of the dynamics leading up to that event. Imagine extending the network in Figure 1 to much larger times, with occasional interactions between the two systems. Suppose we control the mechanisms that initialize and evolve System 2 (the memory), but System 1 always takes random steps. Then, $S_{2,u}$ cannot know the state of System 1 after the most recent interaction time $t < u$. Indeed, by the d-separation criterion, conditional on $S_{1,t}$, any subsequent random steps taken by System 1 are independent of $S_{2,u}$.

Note that our explanation makes no use of the thermodynamic arrow. In contrast, Wolpert and Kipper [11,12] argue that photograph-type memories require an entropy-producing initialization step. While entropy is produced in almost all real-life examples, in some sense this is a trivial observation: a Universe with a thermodynamic arrow will certainly never lose entropy, and it is rather difficult to keep entropy exactly constant. This applies not only to memories, but to all physical processes. Nonetheless, Bennett [65] shows that in principle, memory initialization can be performed in a thermodynamically reversible (entropy-preserving) manner.

Another way to see that the classical second law of thermodynamics (Theorem 1) cannot possibly tell the whole story is to consider the joint entropy of two non-interacting systems (i.e., independent Markov chains)

$$H(S_{t,\{1,2\}}) = H(S_{t,1}) + H(S_{t,2}) - I(S_{t,1} : S_{t,2}).$$

Theorem 1 applies to each of the systems as well as to their union. Hence, the total entropy $H(S_{t,\{1,2\}})$, as well as each of the terms $H(S_{t,1})$ and $H(S_{t,2})$, are non-decreasing.

This does not rule out the possibility of the systems spontaneously becoming correlated: the second law is consistent with an increase in $I(S_{t,1} : S_{t,2})$, provided that the other terms increase as well to compensate. The data processing inequality ([54] Theorem 2.8.1) forbids this as a separate consequence, not of Theorem 1, but of the causal assumptions.

It therefore seems reasonable to think of the arrows of time as fundamentally originating from causality, rather than from the second law of thermodynamics.

4.4. The Evolutionary Arrow of Time

Darwinian evolution is a process by which creatures adapt to their environment, as a result of repeated exposure to its selective pressures over the span of many generations. While environments can change, the homogeneity of the laws of physics across space and time provide some degree of consistency.

As we saw previously, for a causal multibaker map with a time-independent law f_t , the coarse-grained forward-time probabilities are also time-independent, while the reverse-time probabilities generally are not. If we try to cheat the arrow of time by hardcoding advanced creatures into the initial condition, we must ask: what are these creatures adapted to? They cannot have adapted to the underlying law f_t , because the coarse-grained reverse-time trajectory would always converge to the initial condition, regardless of f_t .

In contrast, we saw that the coarse-grained *forward*-time probabilities are locally determined by f_t . They determine the statistical properties of emergent processes, such as genetic mutation, survival, and reproduction, which in turn determine the natural selection gradient. If the law f_t is applied consistently across a large environment over many generations, then the gene pool evolves along that gradient.

4.5. The Agential/Volitional Arrow of Time

If causal multibaker maps could support the evolution of intelligent agents, then perhaps we should not be so surprised if those agents come to regard the past as “fixed”, and the future as “open” to influence from their actions. When discussing the epistemic arrow, we saw that there exist mechanisms which make reliable records of past events; no action from an agent can invalidate a record, short of vandalizing the record itself. In contrast, we saw that no such records exist for future events.

The evolutionary arrow yields additional insights. In the context of Darwinian evolution, the interventions and counterfactuals corresponding to our alternative actions take on a very real meaning: they model competing agents who respond to the same situation with different algorithms. Some of these agents will be more successful than others. An intelligence with the ability to imagine itself in the role of every possible agent, and then acting as the most successful one, will enjoy the maximum survival advantage. Therefore, natural selection increases the prevalence of intelligences that behave this way.

Rovelli [4,18] presents a different account of the agential arrow, in which decisions are considered to be random and entropy-producing. Our approach is closer to that of Rehn [60]: we do not require decisions to be random, but instead view them as outputs of an algorithm that computes estimated outcomes for many different actions.

5. Discussion

Historically, the lack of a mathematically precise model for emergent time-reversal asymmetries posed a major obstacle to their detailed study. While the original multibaker maps provided useful models for thermodynamics [39,40], they lacked the causal interactions responsible for the epistemic, evolutionary, and agential arrows of time. Meanwhile, Pearlean causal models successfully captured these asymmetries [9], but provided no connection to reversible dynamical laws.

Bridging these ideas together, our causal multibaker maps coarse-grain into persistent causal models, offering a tractable, precise framework that can be configured for arbitrary causal interactions in discrete time and space. Follow-up work can split in two main directions.

The first is to study the properties of persistent causal models and use them to model additional phenomena. In particular, the evolutionary and agential arrows should be investigated in much greater depth. Papadopoulos et al. [66] find that machine prediction performs better in a forward-time direction, suggesting that learning algorithms might take advantage of the causal structure (1). Causal modeling may also yield insights on the thermodynamics of biochemical and computational systems [47,67].

The second direction is to investigate how the emergence of time asymmetries in the real Universe resembles or differs from their emergence in causal multibaker maps. Our maps are highly stylized, ignoring many important aspects of real physics. For example, spacetime has a manifold structure, whose most obvious asymmetry is the fact that it is expanding; this is the famous **cosmological arrow of time** [68]. Quantum extensions of causal modeling concepts are another highly active area of research, that will likely require reevaluating our classical intuitions [69–75]. Since the discrete analogue of a field theory is a lattice or cellular automaton, it would also be worth exploring how causal sparsity arises in such models [33,52,76].

A more subtle difference is that causal multibaker maps ensure the Markov property by conveniently shifting their microvariables r_k . In the real Universe, it is much less clear whether and how the Markov property arises [34–38]. In place of our microvariables, chaos or quantum decoherence might play a role in providing independent sources of randomness [77–79].

The epistemic arrow exposes a deeper issue: records only give us information if we know how to interpret them. To do so, we require prior knowledge of the memory's mechanism and initial state. If we feign ignorance and assume a uniform Bayesian prior on the state of the Universe, then we effectively treat the Universe as if it were at heat death. By the second law of thermodynamics, we can never return to a state of non-ignorance, as every observation would be suspected of being a random “Boltzmann brain” fluctuation [33,80,81]. As explained by Wolpert and Kipper [12], this issue is closely related to formal impossibility results in the theory of inductive inference [82,83].

The question becomes: how should a living being embedded inside a causal multibaker map infer non-uniform probabilities such as (2)? Scharnhorst et al. [84] suggest that we should condition on some trustworthy physics knowledge: specifically, the initial and current entropy and dynamics of the Universe. Another approach, based on Occam's razor, uses algorithmic information theory to define a simplicity prior [85–88]. Algorithmic information theory might add more nuance to our understanding of the arrows of time, as it defines probability-free notions of entropy and causality [89–92], as well as alternative measures of complexity such as Bennett's logical depth [93,94].

Finally, we argued that evolution favors agents who reason in terms of forward-time causal interventions. However, in settings where multiple agents must reason about each other's behavior, Yudkowsky, Soares, and Levinstein [95,96] argue that it is sometimes advantageous to choose actions as if they could affect the physical past. In light of ongoing advances in the psychology and neuroscience of time perception [97–100], one could try to study whether human psychology includes such exceptions. As human and artificial intelligence continues to advance and collaborate at larger scales, it becomes increasingly important to clarify the causal foundations of inference and decision theory.

Author Contributions: Conceptualization, A.E.; methodology, A.E.; formal analysis, A.E.; writing—original draft preparation, A.E.; writing—review and editing, A.E. and M.H.; supervision, M.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The original contributions presented in the study are included in the article, further inquiries can be directed to the corresponding author.

Acknowledgments: The first author is indebted to discussions with Jason Li.

Conflicts of Interest: Author Marcus Hutter was employed by Google LLC. Both authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

- Ehrenfest, P.; Ehrenfest, T. *The Conceptual Foundations of the Statistical Approach in Mechanics*; Courier Corporation: North Chelmsford, MA, USA, 1959.
- Davies, P.C.W. *The Physics of Time Asymmetry*; University of California Press: Berkeley, CA, USA, 1977.
- Reichenbach, H. *The Direction of Time*; University of California Press: Berkeley, CA, USA, 1956; Volume 65.
- Rovelli, C. How oriented causation is rooted into thermodynamics. *Philos. Phys.* **2023**, *1*, 11. [[CrossRef](#)]
- Bell, J.S. The theory of local beables. *Epistemol. Lett.* **1975**, *9*, 11–24.
- Bell, J.S. Free variables and local causality. *Epistemol. Lett.* **1977**, *15*, 79–84.
- Granger, C.W.J. Some recent development in a concept of causality. *J. Econom.* **1988**, *39*, 199–211. [[CrossRef](#)]
- Imbens, G.W.; Rubin, D.B. *Causal Inference in Statistics, Social, and Biomedical Sciences*; Cambridge University Press: Cambridge, UK, 2015.
- Pearl, J. *Causality*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2009.
- Hernán, M.A.; Robins, J.M. *Causal Inference: What If*; Boca Raton: Chapman & Hall/CRC: Boca Raton, FL, USA, 2020.
- Wolpert, D.H. Memory systems, computation, and the second law of thermodynamics. *Int. J. Theor. Phys.* **1992**, *31*, 743–785. [[CrossRef](#)]
- Wolpert, D.H.; Kipper, J. Memory Systems, the Epistemic Arrow of Time, and the Second Law. *Entropy* **2024**, *26*, 170. [[CrossRef](#)]
- Hawking, S. *The Illustrated A Brief History of Time: Updated and Expanded Edition*; Bantam: New York, NY, USA, 1996.
- Mlodinow, L.; Brun, T.A. Relation between the psychological and thermodynamic arrows of time. *Phys. Rev. E Stat. Nonlinear Soft Matter Phys.* **2014**, *89*, 052102. [[CrossRef](#)]
- Gregersen, N.H. *From Complexity to Life: On the Emergence of Life and Meaning*; Oxford University Press: Oxford, UK, 2002.
- Blum, H.F. *Time's Arrow and Evolution*; Princeton University Press: Princeton, NJ, USA, 2015; Volume 2075.
- Donnan, F.G. Activities of life and the second law of thermodynamics. *Nature* **1934**, *133*, 99. [[CrossRef](#)]
- Rovelli, C. The thermodynamic cost of choosing. *Found. Phys.* **2024**, *54*, 28. [[CrossRef](#)]
- Gibbard, A.; Harper, W.L. Counterfactuals and two kinds of expected utility. In *Ifs*; Springer: Berlin/Heidelberg, Germany, 1978; pp. 153–190.
- Lewis, D. Counterfactual dependence and time's arrow. *Noûs* **1979**, *13*, 455–476. [[CrossRef](#)]
- Ringström, H. *The Cauchy Problem in General Relativity*; European Mathematical Society: Helsinki, Finland, 2009; Volume 6.
- Lewis, D. Are we free to break the laws? *Theoria* **1981**, *47*, 113–121. [[CrossRef](#)]
- Jost, R. A remark on the C.T.P. theorem. *Helv. Phys. Acta* **1957**, *30*, 409–416.
- Albert, D.Z. *Time and Chance*; American Association of Physics Teachers: College Park, MD, USA, 2001.
- Loewer, B. The mentaculus vision. In *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature*; World Scientific: Singapore, 2020; pp. 3–29.
- Seifert, U. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* **2012**, *75*, 126001. [[CrossRef](#)]
- Peliti, L.; Pigolotti, S. *Stochastic Thermodynamics: An Introduction*; Princeton University Press: Princeton, NJ, USA, 2021.
- Gaspard, P. *The Statistical Mechanics of Irreversible Phenomena*; Cambridge University Press: Cambridge, UK, 2022.
- Shiraishi, N. *An Introduction to Stochastic Thermodynamics: From Basic to Advanced*; Springer Nature: Berlin/Heidelberg, Germany, 2023; Volume 212.
- Cover, T.M. Which processes satisfy the second law. In *Physical Origins of Time Asymmetry*; Cambridge University Press: Cambridge, UK, 1994; pp. 98–107.
- Janzing, D. On the entropy production of time series with unidirectional linearity. *J. Stat. Phys.* **2010**, *138*, 767–779. [[CrossRef](#)]
- Janzing, D. The cause-effect problem: Motivation, ideas, and popular misconceptions. In *Cause Effect Pairs in Machine Learning*; The Springer Series on Challenges in Machine Learning; Springer: Cham, Switzerland, 2019; pp. 3–26. [[CrossRef](#)]
- Ebtekar, A. Information dynamics and the arrow of time. *arXiv* **2021**, arXiv:2109.09709.
- Nicolis, G.; Nicolis, C. Master-equation approach to deterministic chaos. *Phys. Rev. A Gen. Phys.* **1988**, *38*, 427–433. [[CrossRef](#)]
- Werndl, C. Are deterministic descriptions and indeterministic descriptions observationally equivalent? *Stud. Hist. Philos. Sci. Part B Stud. Hist. Philos. Mod. Phys.* **2009**, *40*, 232–242. [[CrossRef](#)]
- Figueroa-Romero, P.; Modi, K.; Pollock, F.A. Almost Markovian processes from closed dynamics. *Quantum* **2019**, *3*, 136. [[CrossRef](#)]
- Figueroa-Romero, P.; Pollock, F.A.; Modi, K. Markovianization with approximate unitary designs. *Commun. Phys.* **2021**, *4*, 127. [[CrossRef](#)]
- Strasberg, P.; Winter, A.; Gemmer, J.; Wang, J. Classicality, Markovianity, and local detailed balance from pure-state dynamics. *Phys. Rev. A* **2023**, *108*, 012225. [[CrossRef](#)]
- Gaspard, P. Diffusion, effusion, and chaotic scattering: An exactly solvable Liouvillian dynamics. *J. Stat. Phys.* **1992**, *68*, 673–747. [[CrossRef](#)]
- Altaner, B.; Vollmer, J. A microscopic perspective on stochastic thermodynamics. *arXiv* **2012**, arXiv:1212.4728.

41. Neyman, J. Sur les applications de la théorie des probabilités aux expériences agricoles: Essai des principes. *Rocz. Nauk. Rol.* **1923**, *10*, 1–51.
42. Lewis, D. Causation. *J. Philos.* **1973**, *70*, 556–567. [[CrossRef](#)]
43. Granger, C.W.J. Investigating causal relations by econometric models and cross-spectral methods. *Econom. J. Econom. Soc.* **1969**, *37*, 424–438. [[CrossRef](#)]
44. Holland, P.W. Statistics and causal inference. *J. Am. Stat. Assoc.* **1986**, *81*, 945–960. [[CrossRef](#)]
45. Spirtes, P.; Glymour, C.; Scheines, R. *Causation, Prediction, and Search*; MIT Press: Cambridge, MA, USA, 2001.
46. Ito, S.; Sagawa, T. Information thermodynamics on causal networks. *Phys. Rev. Lett.* **2013**, *111*, 180603. [[CrossRef](#)] [[PubMed](#)]
47. Ito, S. *Information Thermodynamics on Causal Networks and Its Application to Biochemical Signal Transduction*; Springer: Berlin/Heidelberg, Germany, 2016.
48. Wolpert, D.H. Uncertainty relations and fluctuation theorems for Bayes nets. *Phys. Rev. Lett.* **2020**, *125*, 200602. [[CrossRef](#)]
49. Carroll, S.M. The Arrow of Time in Causal Networks. Invited presentation at Simons Institute Causality Program, 22 April 2022. Available online: <https://www.youtube.com/watch?v=6slug9rjalQ> (accessed on 26 May 2024).
50. Lind, D.; Marcus, B. *An Introduction to Symbolic Dynamics and Coding*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2021.
51. Devaney, R.L. *An Introduction To Chaotic Dynamical Systems*, 3rd ed.; CRC Press: Boca Raton, FL, USA, 2021.
52. Kaneko, K. Overview of coupled map lattices. *Chaos Interdiscip. J. Nonlinear Sci.* **1992**, *2*, 279–282. [[CrossRef](#)]
53. Vichniac, G.Y. Simulating physics with cellular automata. *Phys. D Nonlinear Phenom.* **1984**, *10*, 96–116. [[CrossRef](#)]
54. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, 2nd ed.; Wiley-Interscience: Hoboken, NJ, USA, 2006.
55. Dondelinger, F.; Lèbre, S.; Husmeier, D. Non-homogeneous dynamic Bayesian networks with Bayesian regularization for inferring gene regulatory networks with gradually time-varying structure. *Mach. Learn.* **2013**, *90*, 191–230. [[CrossRef](#)]
56. Robinson, J.; Hartemink, J. Non-stationary dynamic Bayesian networks. *Adv. Neural Inf. Process. Syst.* **2008**, *21*, 1369–1376.
57. Song, L.; Kolar, M.; Xing, E. Time-varying dynamic Bayesian networks. *Adv. Neural Inf. Process. Syst.* **2009**, *22*, 1732–1740.
58. Verma, T.S.; Pearl, J. Causal networks: Semantics and expressiveness. In Proceedings of the Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN, USA, 19–21 August 1988; pp. 352–359.
59. Pearl, J. Structural Counterfactuals: A Brief Introduction. *Cogn. Sci.* **2013**, *37*, 977–985. [[CrossRef](#)] [[PubMed](#)]
60. Rehn, E.M. Free Will Belief as a Consequence of Model-Based Reinforcement Learning. In Proceedings of the International Conference on Artificial General Intelligence, Seattle, WA, USA, 19–22 August 2022; Springer: Berlin/Heidelberg, Germany, 2022; pp. 353–363.
61. Révész, P. A probabilistic solution of problem 111. of G. Birkhoff. *Acta Math. Acad. Sci. Hung.* **1962**, *13*, 187–198. [[CrossRef](#)]
62. Boyle, L.; Finn, K.; Turok, N. CPT-symmetric universe. *Phys. Rev. Lett.* **2018**, *121*, 251301. [[CrossRef](#)]
63. Boyle, L.; Finn, K.; Turok, N. The Big Bang, CPT, and neutrino dark matter. *Ann. Phys.* **2022**, *438*, 168767.
64. Maruyama, K.; Nori, F.; Vedral, V. Colloquium: The physics of Maxwell’s demon and information. *Rev. Mod. Phys.* **2009**, *81*, 1. [[CrossRef](#)]
65. Bennett, C.H. The thermodynamics of computation—A review. *Int. J. Theor. Phys.* **1982**, *21*, 905–940.
66. Papadopoulos, V.; Jérémie, W.; Hongler, C. Arrows of Time for Large Language Models. In Proceedings of the Forty-First International Conference on Machine Learning, Vienna, Austria, 21–27 July 2024.
67. Wolpert, D.; Korbel, J.; Lynn, C.; Tasnim, F.; Grochow, J.; Kardeş, G.; Aimone, J.; Balasubramanian, V.; de Giuli, E.; Doty, D.; et al. Is stochastic thermodynamics the key to understanding the energy costs of computation? *arXiv* **2023**, arXiv:2311.17166.
68. Hawking, S.W. Arrow of time in cosmology. *Phys. Rev. D* **1985**, *32*, 2489. [[CrossRef](#)] [[PubMed](#)]
69. Hardy, L. Towards quantum gravity: A framework for probabilistic theories with non-fixed causal structure. *J. Phys. A Math. Theor.* **2007**, *40*, 3081. [[CrossRef](#)]
70. Chiribella, G.; D’Ariano, G.M.; Perinotti, P.; Valiron, B. Quantum computations without definite causal structure. *Phys. Rev. A—At. Mol. Opt. Phys.* **2013**, *88*, 022318. [[CrossRef](#)]
71. Costa, F.; Shrapnel, S. Quantum causal modelling. *New J. Phys.* **2016**, *18*, 063032. [[CrossRef](#)]
72. Allen, J.M.A.; Barrett, J.; Horsman, D.C.; Lee, C.M.; Spekkens, R.W. Quantum common causes and quantum causal models. *Phys. Rev. X* **2017**, *7*, 031021. [[CrossRef](#)]
73. Barrett, J.; Lorenz, R.; Oreshkov, O. Quantum causal models. *arXiv* **2019**, arXiv:1906.10726.
74. Pienaar, J. Quantum causal models via quantum Bayesianism. *Phys. Rev. A* **2020**, *101*, 012104. [[CrossRef](#)]
75. Lorenz, R. Quantum causal models: The merits of the spirit of Reichenbach’s principle for understanding quantum causal structure. *Synthese* **2022**, *200*, 424. [[CrossRef](#)]
76. Kari, J. Reversible Cellular Automata: From Fundamental Classical Results to Recent Developments. *New Gener. Comput.* **2018**, *36*, 145–172. [[CrossRef](#)]
77. Peres, A. Stability of quantum motion in chaotic and regular systems. *Phys. Rev. A* **1984**, *30*, 1610. [[CrossRef](#)]
78. Zurek, W.H. Decoherence, chaos, quantum-classical correspondence, and the algorithmic arrow of time. *Phys. Scr.* **1998**, *1998*, 186. [[CrossRef](#)]
79. Esposito, M.; Lindenberg, K.; den Broeck, C.V. Entropy production as correlation between system and reservoir. *New J. Phys.* **2010**, *12*, 013013. [[CrossRef](#)]
80. Barrow, J.D.; Tipler, F.J.; Anderson, J.L. *The Anthropic Cosmological Principle*; Oxford University Press: New York, NY, USA, 1987.

81. Carroll, S.M. Why Boltzmann brains are bad. In *Current Controversies in Philosophy of Science*; Routledge: Abingdon, UK, 2020; pp. 7–20.
82. Adam, S.P.; Alexandropoulos, S.A.N.; Pardalos, P.M.; Vrahatis, M.N. No free lunch theorem: A review. In *Approximation and Optimization: Algorithms, Complexity and Applications*; Springer: Cham, Switzerland, 2019; pp. 57–82.
83. Wolpert, D.H. The implications of the no-free-lunch theorems for meta-induction. *J. Gen. Philos. Sci.* **2023**, *54*, 421–432. [[CrossRef](#)]
84. Scharnhorst, J.; Wolpert, D.; Rovelli, C. Boltzmann Bridges. *arXiv* **2024**, arXiv:2407.02840.
85. Solomonoff, R.J. A formal theory of inductive inference. Part I. *Inf. Control* **1964**, *7*, 1–22. [[CrossRef](#)]
86. Solomonoff, R.J. A formal theory of inductive inference. Part II. *Inf. Control* **1964**, *7*, 224–254. [[CrossRef](#)]
87. Rathmanner, S.; Hutter, M. A Philosophical Treatise of Universal Induction. *Entropy* **2011**, *13*, 1076–1136. [[CrossRef](#)]
88. Müller, M.P. Law without law: From observer states to physics via algorithmic information theory. *Quantum* **2020**, *4*, 301. [[CrossRef](#)]
89. Zurek, W.H. Algorithmic randomness and physical entropy. *Phys. Rev. A* **1989**, *40*, 4731. [[CrossRef](#)]
90. Gács, P. The Boltzmann entropy and randomness tests. In Proceedings of the Proceedings Workshop on Physics and Computation, PhysComp'94, Dallas, TX, USA, 17–20 November 1994; pp. 209–216.
91. Janzing, D.; Schölkopf, B. Causal inference using the algorithmic Markov condition. *IEEE Trans. Inf. Theory* **2010**, *56*, 5168–5194. [[CrossRef](#)]
92. Ebtekar, A.; Hutter, M. Foundations of algorithmic thermodynamics. *Manuscript submitted for publication*.
93. Bennett, C.H. *Logical Depth and Physical Complexity*; Oxford University Press: Oxford, UK, 1988; pp. 227–257.
94. Bennett, C.H. Complexity in the universe. In *Physical Origins of Time Asymmetry*; Cambridge University Press: Cambridge, UK, 1994; pp. 33–46.
95. Yudkowsky, E.; Soares, N. Functional Decision Theory: A New Theory of Instrumental Rationality. *arXiv* **2017**, arXiv:1710.05060.
96. Levinstein, B.A.; Soares, N. Cheating death in damascus. *J. Philos.* **2020**, *117*, 237–266. [[CrossRef](#)]
97. Wittmann, M.; Paulus, M.P. Decision making, impulsivity and time perception. *Trends Cogn. Sci.* **2008**, *12*, 7–12. [[CrossRef](#)] [[PubMed](#)]
98. Weger, U.W.; Pratt, J. Time flies like an arrow: Space-time compatibility effects suggest the use of a mental timeline. *Psychon. Bull. Rev.* **2008**, *15*, 426–430. [[CrossRef](#)]
99. Grondin, S. Timing and time perception: A review of recent behavioral and neuroscience findings and theoretical directions. *Atten. Percept. Psychophys.* **2010**, *72*, 561–582. [[CrossRef](#)]
100. Gauthier, B.; Pestke, K.; van Wassenhove, V. Building the arrow of time... over time: A sequence of brain activity mapping imagined events in time and space. *Cereb. Cortex* **2019**, *29*, 4398–4414. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.