Heraclitus-Maximal Worlds*

JB Manchak and Thomas William Barrett

[Forthcoming in the *Journal of Philosophical Logic*]

Abstract

Within the context of general relativity, the Heraclitus asymmetry property requires that no distinct pair of spacetime events have the same local structure (Manchak and Barrett 2023). Here, we explore Heraclitus-maximal worlds – those which are "as large as they can be" with respect to the Heraclitus property. Using Zorn's lemma, we prove that such worlds exist and highlight a number of their properties. If attention is restricted to Heraclitus-maximal worlds, we show a sense in which observers have the epistemic resources to know which world they inhabit.

1 Introduction

Within the context of general relativity, Leibnizian metaphysics seems to demand that worlds are "maximal" with respect to a variety of spacetime properties (Geroch 1970; Earman 1995). Here, we explore maximal worlds with respect to the "Heraclitus" asymmetry property which demands that no distinct pair of spacetime events have the same local structure (Barrett and Manchak 2023). We will show that Heraclitus-maximal worlds exist and that every Heraclitus world is contained in some Heraclitus-maximal world. This amounts to a type of compatibility between the Leibnizian and Heraclitian demands. We then go on to show a sense of incompatibility between these demands and the existence of non-isomorphic but "observationally indistinguishable" worlds introduced by Glymour (1972, 1977) and Malament (1977).

Our discussion in this paper synthesizes three separate literatures: the physics and philosophy of spacetime symmetry and structure (Weyl 1952; Stein 1967; Earman 1989; Dasgupta 2016; Barrett 2018; North 2021), and the modal structure of spacetime, especially with respect to Leibnizian maximality (Penrose 1969; Geroch 1970; Clarke 1993; Earman 1995; Manchak 2016; Krasnikov 2018); the possibility of cosmic underdetermination (Glymour 1977; Malament 1977; Manchak 2009; Norton 2011; Butterfield 2014; Belot 2023). We will see that the three subjects interact in fruitful ways. It is also somewhat remarkable that

*We wish to thank David Malament and two anonymous referees for a number of helpful comments.

almost none of the technical details of general relativity will be needed in what follows. For the most part, proofs come down to basic set theory including simple applications of Zorn's lemma.

2 Sub-worlds and Isomorphisms

One starts with the collection $\mathscr{U}$ of models of (standard) general relativity. This is the collection of smooth, connected, Hausdorff manifolds equipped with a smooth, Lorentzian metric which is assumed to be time-orientable (Hawking and Ellis 1973). In what follows, we shall refer to elements of $\mathscr{U}$ as "worlds" although it is important to recognize that, strictly speaking, such elements are mere mathematical models of worlds – geometric objects – and not worlds themselves.[1] Only two foundational notions will be needed in much of what follows: sub-worlds and isomorphisms. With these notions, one can define Heraclitus-maximal worlds and explore their basic properties. A third foundational notion – the past of an event – is a special type of sub-world needed later on to discuss the epistemic issues.

Consider sub-worlds first. The manifold structure of a world $W \in \mathscr{U}$ includes a manifold topology on W. Any spacetime region $O \subseteq W$ which is both open and connected in the manifold topology on W inherits a manifold and metric structure from W. This O then counts as a world in its own right: $O \in \mathscr{U}$. Call such a region $O \subseteq W$ a **sub-world** of the world W. One can show that sub-worlds of any world $W \in \mathscr{U}$ can be arbitrarily small in the sense that for any open set $S \subseteq W$ however small, there exists a sub-world O of W which fits inside S. Sub-worlds can also be big in the sense that a world W always counts as a sub-world of itself.

Now consider isomorphisms. We say worlds $W_1, W_2 \in \mathscr{U}$ are **isomorphic** if there is a bijection $\psi : W_1 \to W_2$ such that both it and its inverse preserve all manifold and metric structure. In the natural way, isomorphisms between worlds gives rise to an equivalence relation $\sim$ on $\mathscr{U}$; for any worlds $W_1, W_2 \in \mathscr{U}$, the equivalence classes $[W_1], [W_2] \in \mathscr{U}/\sim$ are such that $[W_1] = [W_2]$ if and only if W_1 and W_2 are isomorphic (see Hawking and Ellis 1973, p. 56). In what follows, equivalence classes of isomorphic worlds will prove useful to consider in applying Zorn's lemma (Wald 1984, p. 263; Sbierski 2016, p. 305). But we emphasize that no philosophical assumptions are made or needed here concerning the "equivalence" or "identity" of isomorphic worlds.[2]

For any worlds $W_1, W_2 \in \mathscr{U}$, an **isomorphic embedding** is an injective map $\theta : W_1 \to W_2$ which preserves all manifold and metric structure in the sense that $\theta[W_1]$ is a (not necessarily proper) sub-world of W_2 which is isomorphic to W_1. For any worlds $W_1, W_2 \in \mathscr{U}$, we say W_2 is a (not necessarily proper)

[1] For a recent discussion on the distinction between general relativistic models and the worlds they represent, see Fletcher (2020), Roberts (2020), and Bradley and Weatherall (2020).

[2] Various assumptions of this kind are sometimes called "Leibniz equivalence" in the hole argument literature (see Earman and Norton 1987; Belot 2017; Weatherall 2018; Roberts 2020).

extension of W_1 if there is an isomorphic embedding $\theta : W_1 \to W_2$. Let $\leq$ be the relation on $\mathscr{U}/\!\sim$ such that $[W_1] \leq [W_2]$ if and only if any element of $[W_2]$ is an extension of any element of $[W_1]$. This relation is clearly reflexive and transitive but, somewhat surprisingly, it fails to be a partial order since it is not anti-symmetric: there are worlds $W_1, W_2 \in \mathscr{U}$ such that $[W_1] \leq [W_2]$ and $[W_2] \leq [W_1]$ and yet $[W_1] \neq [W_2]$ (Geroch 1970, p. 276).

To see this, start with a two-dimensional Minkowski world $M \in \mathscr{U}$ in standard (t, x) coordinates (Hawking and Ellis 1973, p. 118). This is the world of special relativity. Now consider the sub-world $M_0 \subset M$ which is the $t > 0$ portion of M. Also consider the sub-world $M_0 - \{e\} \subset M$ where e is the event $(1, 0) \in M_0$. It is immediate that M_0 is an extension of its sub-world $M_0 - \{e\}$; just consider the inclusion map from the latter to the former. Because of the symmetries of the two worlds, the other direction also holds. To see this, consider the map $\theta : M_0 \to M_0 - \{e\}$ defined by $\theta(t, x) = (t + 1, x)$; it is an isomorphic embedding (see Figure 1). So $[M_0] \leq [M_0 - \{e\}]$ and $[M_0 - \{e\}] \leq [M_0]$. But $[M_0] \neq [M_0 - \{e\}]$ since M_0 is not isomorphic to $M_0 - \{e\}$ due to the "missing" event e (the former world is topologically simply connected while the latter is not).

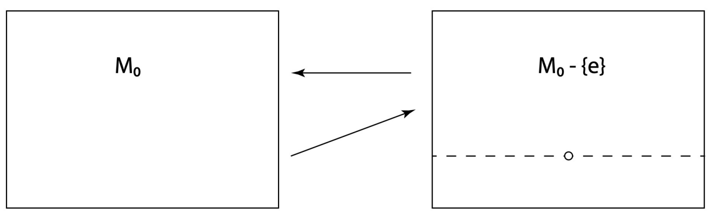

Figure 1: The world $M_0 - \{e\}$ can be isomorphically embedded into M_0 using the inclusion map. The world M_0 can be isomorphically embedded into the region above the dotted line in $M_0 - \{e\}$.

Given that the relation $\leq$ on $\mathscr{U}/\!\sim$ fails to be a partial order, it has become customary to work with the collection of all "framed" worlds when exploring maximality issues (Geroch 1970; Earman 1995; Krasnikov 2018). One frames a world by associating with it an orthonormal basis of vectors at some event. This additional structure requires that isomorphic embeddings that preserve the frame are unique; one can then construct a partial order on the collection of equivalence classes of framed worlds. Zorn's lemma is used to secure the existence of a maximal framed extension for any given framed world. One then "throws away" the frames to deduce the existence of a maximal world for any given world. Employing framed worlds is certainly a useful workaround when it is needed. But as we will see below, it is not necessary to introduce and remove structure in this way when one limits attention to Heraclitus worlds. Within that context, isomorphic embeddings are unique and therefore the relation $\leq$ defined above counts as a partial order on the collection of equivalence classes of Heraclitus worlds.

3 Heraclitus Worlds

One can use the notions of sub-worlds and isomorphisms to define the Heraclitus asymmetry property (Manchak and Barrett 2023).

Definition. A world $W \in \mathscr{U}$ is **Heraclitus** if, for any distinct events $e_1, e_2 \in W$ and any sub-worlds $O_1, O_2 \subseteq W$ containing e_1 and e_2 respectively, there is no isomorphism $\psi : O_1 \to O_2$ such that $\psi(e_1) = e_2$.

In a Heraclitus world, no distinct events have the same local structure – each event is unlike every other. One might say of such a world that one cannot step twice into the same river. Heraclitus worlds are completely devoid of symmetries. Not only does a Heraclitus world fail to have non-trivial global symmetries in the sense that the only isomorphism from it to itself is the identity map, such a world also fails to have non-trivial local symmetries in the sense that no distinct points have isomorphic local neighborhoods. The Heraclitus asymmetry property allows one to prove a variety of uniqueness results. For example, one can show a sense in which, if Heraclitus worlds share the same "local" properties (see the definition in Section 5), they must share all properties – they are isomorphic (Manchak and Barrett 2023). Another related example will be given below concerning observationally indistinguishable worlds.

We know that Heraclitus worlds exist (Manchak and Barrett 2023). Consider three basic properties of Heraclitus worlds which will be used frequently in what follows.

(H1) Any sub-world of a Heraclitus world is a Heraclitus world. This is immediate from the definition of a Heraclitus world.

(H2) Isomorphic embeddings among Heraclitus worlds are unique. To see this, consider isomorphic embeddings $\theta_1, \theta_2 : W_1 \to W_2$ where W_1 and W_2 are Heraclitus worlds. If there were an event $e \in W_1$ such that $\theta_1(e) \neq \theta_2(e)$, then there would be an isomorphism $\psi : \theta_1[W_1] \to \theta_2[W_1]$ given by $\psi = \theta_2 \circ \theta_1^{-1}$ such that $\psi(\theta_1(e)) = \theta_2(e)$. So W_2 would fail to be Heraclitus: a contradiction.

(H3) A Heraclitus world cannot properly extend itself (i.e. any world isomorphic to it). To see this, suppose there are isomorphic Heraclitus world W_1 and W_2 such that W_2 properly extends W_1. So there must be distinct isomorphic embeddings from W_1 to W_2: one proper and one not. But this contradicts (H2).

Consider the collection $\mathscr{H} \subset \mathscr{U}$ of Heraclitus worlds. We find that the relation $\leq$ defined above on $\mathscr{U}/\sim$ counts as a partial order when restricted to $\mathscr{H}/\sim$. Reflexivity and transitivity are inherited from the relation $\leq$ on $\mathscr{U}/\sim$. We need only show antisymmetry. Suppose $[W_1] \leq [W_2]$ and $[W_2] \leq [W_1]$ for some worlds $W_1, W_2 \in \mathscr{H}$. If $[W_2] \neq [W_1]$, then W_1 and W_2 are not isomorphic. In that case, then since $[W_1] \leq [W_2]$ and $[W_2] \leq [W_1]$, each world is a proper extension of the other. So each world counts as a proper extension of itself which is impossible given (H3). So $[W_1] = [W_2]$ and $\leq$ is a partial order on $\mathscr{H}/\sim$. We note here that considering equivalence classes is crucial in establishing the antisymmetry needed for the partial order. Any world $W_1 \in \mathscr{U}$ is isomorphic to some non-identical $W_2 \in [W_1]$. So W_1 can be isomorphically embedded into W_2 and vice versa even though $W_1 \neq W_2$.

Since $\leq$ is a partial order on $\mathscr{H}/\sim$, we will eventually be able to use Zorn's lemma to show the existence of a Heraclitus world which is maximal with respect to that property. To do this, we first need the following.

Lemma 1. Any sub-collection of $\mathscr{H}/\sim$ which is totally ordered by $\leq$ has an upper bound in $\mathscr{H}/\sim$.

Proof. Let $\{X_i\}$ be a collection of equivalence classes of worlds in $\mathscr{H}/\sim$ which is totally ordered by $\leq$. Associate for each equivalence class X_i a representative world $W_i \in X_i$ using the axiom of choice. Whenever $i \leq j$, we have an isomorphic embedding $\theta_{ij} : W_i \to W_j$. From (H2), the embedding is unique. Following Hawking and Ellis (1973, p. 249), construct a world W by taking the "union" of all of the worlds W_i such that whenever $i \leq j$, each event $e \in W_i$ is identified with the event $\theta_{ij}(e) \in W_j$.[3] The manifold and metric structure on the world W are induced from all of the worlds W_i. Since $[W_i] \leq [W]$ for all i, we are done once we verify that $W \in \mathscr{H}$.

Suppose $W \notin \mathscr{H}$. So there are sub-worlds $O_1, O_2 \subseteq W$ containing distinct events e_1, e_2 respectively and an isomorphism $\psi : O_1 \to O_2$ such that $\psi(e_1) = e_2$. We know there is some world W_i such that $e_1, e_2 \in W_i$. Let $U_1 = O_1 \cap W_i$ and let $U_2 = \psi[U_1]$. Although the sub-world U_1 is contained in W_i, the sub-world U_2 may not be. So consider the sub-world $V_2 = U_2 \cap W_i$ and the sub-world $V_1 = \psi^{-1}[V_2]$. Now it follows that $V_1, V_2 \subseteq W_i$. The isomorphism defined by restricting the domain of ψ to V_1 maps the event e_1 to e_2 which contradicts the Heraclitus property of W_i. So $W \in \mathscr{H}$. □

One naturally wonders: how "physically reasonable" are Heraclitus worlds? As we shall see, there exist two-dimensional Heraclitus worlds with well-behaved local and causal structure in the sense that they are globally hyperbolic, vacuum solutions of Einstein's equation which therefore necessarily satisfy all energy conditions (see the discussion after Proposition 5). But it is unknown if there are such worlds that are four-dimensional (see Manchak and Barrett 2023). Because of their asymmetries, Heraclitus worlds are difficult to construct. Even so, such worlds may be turn out to be "generic" among all worlds in $\mathscr{U}$. Indeed, this is the case for worlds with Riemannian geometry (Sunada 1985, Hebda 2010).

In what follows, let $\mathscr{R} \subseteq \mathscr{U}$ represent a collection of "physically reasonable" worlds. One can naturally consider the following asymmetry condition on such

[3] Although Hawking and Ellis (1973) consider the "union" of worlds, strictly speaking, a direct limit is taken. See Geroch (1970) and Wong (2013) for details.

a collection.

(Heraclitus) $\mathscr{R} \subseteq \mathscr{H}$.

Of course, the (Heraclitus) condition may fail – say, if the collections $\mathscr{R}$ and $\mathscr{H}$ are disjoint. In this paper, we simply explore the logical consequences if (Heraclitus) were to hold.

4 Heraclitus-Maximal Worlds

Let us now turn to the task of defining a collection of worlds that are maximal with respect to the Heraclitus asymmetry property. Once again, we need only consider the notions of sub-worlds and isomorphisms to do this. In the natural way, any property of worlds gives rise to a collection $\mathscr{P} \subseteq \mathscr{U}$ of worlds with the property. Following Geroch (1970), we will identify properties of worlds with their associated sub-collections. We are now in a position to formulate the notion of maximality with respect to arbitrary properties.

Definition. For any property $\mathscr{P} \subseteq \mathscr{U}$, a world in $\mathscr{P}$ is **$\mathscr{P}$-maximal** if it has no proper extensions in $\mathscr{P}$.

A $\mathscr{P}$-maximal world is one that is "as large as it can be" with respect to the property $\mathscr{P} \subseteq \mathscr{U}$. Leibnizian metaphysics seems to demand that worlds be maximal with respect to a variety of properties. Consider, for example, any world $W \in \mathscr{U}$ and any spacetime event $e \in W$. The region $W - \{e\}$ counts as a proper sub-world of W. So the world $W - \{e\}$ fails to be $\mathscr{U}$-maximal and would therefore seem to be metaphysically unacceptable: "for the Creative Force to actualize a proper subpart of a larger spacetime would seem to be a violation of Leibniz's principles of sufficient reason and plenitude" (Earman 1995, p. 32). Following this line, it has become dogma that any physically reasonable world must be a $\mathscr{U}$-maximal world (Clarke 1993). But this dogma is empirically unverifiable (see Proposition 4 below) and it also has the potential to clash with other properties of interest. Suppose, for example, that one is also committed to the position that all reasonable worlds must have property $\mathscr{P} \subset \mathscr{U}$. One has no a priori assurance that $\mathscr{P}$-maximality implies $\mathscr{U}$-maximality. Indeed, for a variety of properties of interest, one can find worlds in $\mathscr{P}$ which can be properly extended but only by the "unreasonable" worlds in $\mathscr{U} - \mathscr{P}$ (Manchak 2021). It would therefore be a mistake to consider such worlds physically unreasonable on maximality grounds as the dogma requires.

Given the situation, it seems appropriate to explore Leibnizian maximality not just with respect to the collection of worlds $\mathscr{U}$ but also with respect to a variety of sub-collections $\mathscr{P} \subseteq \mathscr{U}$. A revised form of the dogma requires that a collection of physically reasonable worlds $\mathscr{R} \subseteq \mathscr{U}$ be such that no world in $\mathscr{R}$ can be properly extended by a world in $\mathscr{R}$. This revised dogma can certainly be called into question (see Proposition 4). Indeed, it is unclear whether $\mathscr{R}$-

maximality is a stable property of worlds (Manchak 2023). But it proves quite useful to explore the strength of this revised dogma; consider the following maximality condition.

(Leibniz) Any world in $\mathscr{R}$ is $\mathscr{R}$-maximal.

In what follows, we will be interested in exploring collections $\mathscr{R} \subseteq \mathscr{U}$ that are physically reasonable under the supposition that they satisfy both (Heraclitus) and (Leibniz). We see immediately that if $\mathscr{R} = \mathscr{H}$, then (Heraclitus) is satisfied but (Leibniz) is not. This follows since $\mathscr{H}$ contains proper sub-worlds of Heraclitus worlds which, from (H1), are Heraclitus worlds themselves. Such sub-worlds are worlds in $\mathscr{H}$ that fail to be $\mathscr{H}$-maximal. Now consider the collection $\mathscr{H}^* \subset \mathscr{H}$ of $\mathscr{H}$-maximal worlds. Such worlds cannot be properly extended while still retaining the Heraclitus property. We shall see below that $\mathscr{H}^*$ is non-empty. Here, we verify that both (Heraclitus) and (Leibniz) are satisfied if $\mathscr{R} = \mathscr{H}^*$. Showing the former is immediate since $\mathscr{H}^* \subset \mathscr{H}$. Consider the latter. Let W_1 be any world in $\mathscr{H}^*$ and suppose it fails to be $\mathscr{H}^*$-maximal. So W_1 has a proper extension $W_2 \in \mathscr{H}^* \subset \mathscr{H}$. Since W_2 is Heraclitus and properly extends W_1, it follows that $W_1 \notin \mathscr{H}^*$: a contradiction. So any world in $\mathscr{H}^*$ is $\mathscr{H}^*$-maximal, i.e. any Heraclitus-maximal world is also a (Heraclitus-maximal)-maximal world. So (Leibniz) is also satisfied by $\mathscr{R} = \mathscr{H}^*$.

Aside from $\mathscr{H}^*$, there are a number of other non-empty collections $\mathscr{R} \subseteq \mathscr{U}$ that also satisfy both (Heraclitus) and (Leibniz). It is not difficult to verify that any collection $\mathscr{R} \subseteq \mathscr{H}^*$ will do. But one can also find example collections which are disjoint from $\mathscr{H}^*$. Consider any proper sub-world W of any Heraclitus world and let $\mathscr{R} = \{W\}$. From (H1), W is a Heraclitus world. So (Heraclitus) is true for the property $\{W\}$. From (H3), W cannot properly extend itself. So (Leibniz) is also true for $\{W\}$ even though W is not a Heraclitus-maximal world.

We now show that $\mathscr{H}^*$ is non-empty. Since we know that Heraclitus worlds exist, we are done if we can show the following general result: any Heraclitus world has a Heraclitus-maximal extension. The analogous statement with respect to the collection $\mathscr{U}$ has long been used to underpin the Leibnizian dogma mentioned above; using Zorn's lemma, one can show that any world in $\mathscr{U}$ has a $\mathscr{U}$-maximal extension (Geroch 1970).[4] One might be tempted to conclude that similar results can be obtained for any physically reasonable property. But things are not so simple. For some causal properties, Zorn's lemma cannot be applied and things are left unsettled (Low 2012). For other properties, including the "big bang" property, the analogous statement comes out as false (Manchak 2016). Fortunately, Lemma 1 ensures that things are straightforward with respect to the Heraclitus asymmetry property; using Zorn's lemma, we obtain the

[4]It is an open question whether this foundational statement or Proposition 1 below remain true if Zorn's lemma is not invoked. Recently some work has been done to "dezornify" certain results concerning the existence of "maximal Cauchy developments" (Wong 2013; Sbierski 2016). But such results depend crucially on the property of global hyperbolicity which is not assumed to hold here.

following proposition.

Proposition 1. Any world in $\mathscr{H}$ has an $\mathscr{H}$-maximal extension.

Proof. Let W be a world in $\mathscr{H}$. Let $\mathscr{E} \subset \mathscr{H}$ be the collection of extensions of W in $\mathscr{H}$. Since $\mathscr{H}/\!\sim$ is a partial order, so is $\mathscr{E}/\!\sim$. Consider any collection of equivalence classes of worlds in $\mathscr{E}/\!\sim$ which is totally ordered by $\leq$. By Lemma 1, this collection has an upper bound $X \in \mathscr{H}/\!\sim$. Since $[W] \leq X$, we know that any world in X counts as an extension of W. So $X \in \mathscr{E}/\!\sim$. By Zorn's lemma, there is a maximal element $X^* \in \mathscr{E}/\!\sim$. So for any world $W^* \in X^*$ we see that W^* is an $\mathscr{H}$-maximal extension of W. □

Corollary 1. An $\mathscr{H}$-maximal world exists.

5 Can One Know One's World?

The metric structure of a world $W \in \mathscr{U}$ includes its causal structure. One can use the causal structure to define, for each event $e \in W$, a sub-world $P(e) \subseteq W$ called the (timelike) **past** of e (Hawking and Ellis 1973, p. 183). Under the basic assumption that events outside of $P(e)$ cannot be empirically observed from e, we are able to define a notion of observationally indistinguishable worlds (Malament 1977).

Definition. A world $W_1 \in \mathscr{U}$ is **observationally indistinguishable** from a world $W_2 \in \mathscr{U}$ if, for each event $e_1 \in W_1$, there is an event $e_2 \in W_2$ such that the pasts $P(e_1)$ and $P(e_2)$ are isomorphic.

The definition seems to be a "straightforward rendering of conditions under which observers could not determine the spatio-temporal structure of the universe" (Malament 1977, p. 69). We note here that, although the conditions specified in the definition seem to be sufficient for a type of cosmic underdetermination, they do not seem to be necessary. We will return to this point in the next section.

Note that the relation of observational indistinguishability among worlds is reflexive and transitive but fails to be symmetric: if a world W_1 is observationally indistinguishable from a world W_2, it does not necessarily follow that W_2 is observationally indistinguishable from W_1. Of course, one could consider a stronger symmetric version of observational indistinguishability which does count as an equivalence relation among worlds. An even stronger symmetric version is defined using the pasts of all observer world-lines instead the pasts of all events (Glymour 1972, 1977). We consider the weaker of the three notions here in part because it is still strong enough to signal an underdetermination problem for observers. Moreover, the weaker definition will also allow us to prove the strongest possible uniqueness result below concerning Heraclitus-maximal worlds.

The notion of observational indistinguishability defined gives rise to a general epistemic predicament: modulo modest assumptions, any world is observationally indistinguishable from some other (non-isomorphic) world. To state this result, a few definitions will be useful. Let us say that worlds $W_1, W_2 \in \mathscr{U}$ are **locally isomorphic** if, for each event $e \in W_1$, there is a sub-world O_1 in W_1 containing e that is isomorphic to some sub-world O_2 in W_2, and, correspondingly, with the roles of W_1 and W_2 interchanged. Now say that a property of worlds $\mathscr{P} \subseteq \mathscr{U}$ is **local** if, for any locally isomorphic worlds $W_1, W_2 \in \mathscr{U}$, we have: $W_1 \in \mathscr{P}$ if and only if $W_2 \in \mathscr{P}$. On this definition, one can verify both the Heraclitus property and $\mathscr{U}$-maximality come out as global (non-local) properties. Now, let us say that a world $W \in \mathscr{U}$ has a **God point** if there is an event $e \in W$ such that $P(e) = W$. If a world has a God point, then from that event all events are empirically accessible; an observer can, in principle, see the entire world. Any world W with a God point e permits a type of "time travel" since all events – even those in the future of e – must be in the past of e. We can now formulate a general underdetermination result (Manchak 2009).

Proposition 2. If $W_1 \in \mathscr{U}$ does not have a God point, then there is a non-isomorphic world $W_2 \in \mathscr{U}$ such that W_1 is observationally indistinguishable from W_2 and the worlds share all local properties.

The proposition shows a sense in which (unless one's world has quite strange causal structure) one cannot know which world one inhabits. Moreover, the epistemic predicament persists even if one fixes all local properties. The result has been discussed in a variety of philosophical contexts including those related to induction, scientific realism, and determinism (Norton 2011; Butterfield 2014; Belot 2023). One common response to the result calls into question the physical significance of worlds built in the proof (Cinti and Fano 2021); indeed, a "stupendous 'cut-and-paste' construction" (Butterfield 2012, p. 59) is used whereby an infinite number of copies of the original world are strung together like a clothesline to form the observationally indistinguishable counterpart worlds (cf. Malament 1977). As such, these worlds would seem to be "irrelevant monstrosities by the standards of working cosmologists" (Belot 2023, p. 147).

One principled way to rule out the clothesline construction would be to identify various physically reasonable collections of worlds such that, when attention is restricted to these collections, no epistemic predicament arises. We know this can be done if stronger versions of observational indistinguishability are employed and attention is restricted to certain highly symmetric worlds (Glymour 1977; Belot 2023). Here we show that, even under a weak notion of observational indistinguishability, a general uniqueness result can still be realized for any collection $\mathscr{R} \subseteq \mathscr{U}$ whatsoever that satisfies the (Heraclitus) and (Leibniz) conditions.

Proposition 3. Let $\mathscr{R} \subseteq \mathscr{U}$ be any collection satisfying (Heraclitus) and (Leibniz). For all $W_1, W_2 \in \mathscr{R}$, we have: W_1 is observationally indistinguishable from W_2 if and only if W_1 is isomorphic to W_2.

Proof. Let $\mathscr{R} \subseteq \mathscr{U}$ be any collection satisfying (Heraclitus) and (Leibniz) and consider any worlds $W, W' \in \mathscr{R}$. It follows immediately from the definitions that if W and W' are isomorphic, then each world must be observationally indistinguishable from the other. Suppose W is observationally indistinguishable from W'. Since $\mathscr{R}$ satisfies (Heraclitus), we know that W and W' are Heraclitus. From (H1), we know any sub-worlds of W and W' are also Heraclitus. From (H2), we know that any isomorphic embeddings among such worlds are unique. We will exploit this property often in what follows.

For each $p \in W$, use the axiom of choice to fix a $q \in W$ such that $p \in P(q)$ and define O_p to be the sub-world $P(q)$. Since W is observationally indistinguishable from W', for each event $p \in W$, we can fix an isomorphism $\psi_p : O_p \to O'_p$ where O'_p is some sub-world of W'. From (H2), we know that for any $p, q \in W$, we have $\psi_p = \psi_q$ on the region $O_p \cap O_q$. It follows that the unique map $\psi : \bigcup O_p \to W'$ defined such that $\psi_{|O_p} = \psi_p$ for all $p \in W$ must be smooth (see O'Neill 1983, p. 5). Since $W = \bigcup O_p$, we can restrict the range on ψ to construct a smooth surjective map $\psi' : W \to \psi[W]$.

Next we show that ψ' is injective and therefore a smooth bijection. Let $p, q \in W$ and suppose that $\psi'(p) = \psi'(q)$. It follows that $\psi_p(p) = \psi_q(q)$ where $\psi_p : O_p \to O'_p$ and $\psi_q : O_q \to O'_q$ are the isomorphisms associated with p and q. So $\psi_p^{-1} : O'_p \to O_p$ and $\psi_q^{-1} : O'_q \to O_q$ are isomorphisms. From (H2), we know that $\psi_p^{-1} = \psi_q^{-1}$ on the region $O'_p \cap O'_q$ which contains $\psi'(p) = \psi'(q)$. So $\psi_p^{-1}(\psi_p(p)) = \psi_q^{-1}(\psi_q(q))$ and therefore $p = q$. So ψ' is injective and therefore a smooth bijection.

Next we show that ψ'^{-1} is smooth. Let $p, q \in W$ and consider the isomorphisms $\psi_p^{-1} : O'_p \to O_p$ and $\psi_q^{-1} : O'_q \to O_q$. From (H2), we know that $\psi_p^{-1} = \psi_q^{-1}$ on the region $O'_p \cap O'_q$. Since ψ' is surjective, it follows that $\bigcup O'_p = \psi[W]$. So $\psi'^{-1} : \bigcup O'_p \to W$ is the unique map defined such that $\psi'^{-1}_{|O'_p} = \psi_p^{-1}$ for all $p \in W$. It follows that ψ'^{-1} must be smooth (see O'Neill 1983, p. 5). So $\psi' : W \to \psi[W]$ is a smooth bijection with a smooth inverse: a diffeomorphism. Since ψ' is constructed using the isomorphisms $\psi_p : O_p \to O'_p$ for all $p \in W$, one can verify that this diffeomorphism ψ' preserves all metric structure and therefore counts as an isomorphism.

Since $\psi' : W \to \psi[W]$ is an isomorphism, it follows that $\psi : W \to W'$ is an isomorphic embedding. So W' is an extension of W. Because $\mathscr{R}$ satisfies (Leibniz), we know W' cannot be proper extension of W since this would imply that W is not $\mathscr{R}$-maximal. So W and W' are isomorphic. □

The proposition provides one with a principled way to respond to the underdetermination result given in Proposition 2. If the worlds constructed in its proof are to be dismissed as "irrelevant monstrosities" by working cosmologists, then such a dismissal can now be justified by appeal to the conditions (Heraclitus) and (Leibniz). Given a Heraclitus world W_1, the former condition rules out the clothesline construction of an observationally indistinguishable counterpart world W_2 which is Heraclitus; one cannot string together an infinite number of copies of W_1 to construct W_2 without introducing symmetries into the world. Because isomorphic embeddings among Heraclitus worlds are unique given (H2), and since the the past $P(e)$ of each event $e \in W_1$ must be isomorphically embedded into the observationally indistinguishable counterpart W_2, one can "glue" all of the isomorphic embeddings together to form a single unique isomorphic embedding of all of world W_1 into world W_2. So the (Heraclitus) condition alone is enough to show that W_2 must be an extension of W_1. The (Leibniz) condition can then be invoked to rule out the possibility that W_2 is a proper extension of W_1. This shows that the two worlds must be isomorphic.

Proposition 3 is quite general in the sense that the epistemic predicament dissolves for any collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ that satisfy (Heraclitus) and (Leibniz). We have already shown the existence of a number of distinct collections of worlds which satisfy these conditions including the non-empty collection $\mathscr{H}^*$ of Heraclitus-maximal worlds. So we have the following instantiation of Proposition 3.

Corollary 2. Heraclitus-maximal worlds (which exist) are observationally indistinguishable if and only if they are isomorphic.

Stepping back, we see that much depends on the conditions of (Heraclitus) and (Leibniz). Indeed, one could interpret Proposition 3 as showing just how strong the conditions are. And these conditions can be called into question via a strengthening of Proposition 2. Consider the following.

Proposition 4. If $W_1 \in \mathscr{U}$ does not have a God point, then there is a non-isomorphic world $W_2 \in \mathscr{U}$ such that (i) W_1 is observationally indistinguishable from W_2, (i) W_1 and W_2 share all local properties, (iii) W_2 fails to be $\mathscr{U}$-maximal, and (iv) W_2 fails to be Heraclitus.

The proof for (i)-(iii) is given in Manchak (2011) via the clothesline construction while (iv) follows easily since, as mentioned above, such a construction introduces symmetries into world W_2 which are inconsistent with the Heraclitus property. Proposition 4 can be interpreted as saying that, not only is it impossible to know what world one inhabits, one cannot even know that one's world is Heraclitus or maximal (under the usual definition) through empirical observations. Moreover, such an epistemic predicament persist even if one fixes all local properties. Given the situation, one may want to be open to the failure of the (Heraclitus) and (Leibniz) conditions. If so, one can understand the work in this section neutrally as amounting to a type of no-go result: a certain type of underdetermination is incompatible with the (Heraclitus) and (Leibniz) conditions. But as we will see in the next section, (Heraclitus) and (Leibniz) do not rule out other types of underdetermination.

6 Varieties of Underdetermination

So far, we have considered a certain type of cosmic underdetermination that can arise relative to a particular definition of observationally indistinguishable worlds. As we have mentioned, stronger notions of observational indistinguishability could have also been considered (Glymour 1972, 1977; Malament 1977). These notions would have resulted in stronger forms of underdetermination also inconsistent with the (Heraclitus) and (Leibniz) conditions. But this does not mean that these two conditions are incompatible with all types of physically significant underdetermination. If a world W has no non-isomorphic observationally indistinguishable counterpart world, then there is a sense in which the collective information that all individuals in W have together is sufficient to determine which world they inhabit, but that determination may be beyond the observational reach of any one individual in W (Malament, private communication). We will use this idea to make precise collective vs. individual notions of underdetermination. We will also keep track of a distinction between universal vs. existential types of underdetermination. We go on to show that while some types of underdetermination are ruled out by the (Heraclitus) and (Leibniz) conditions, others are not.

A collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ could be such that all – or merely some – of its worlds inherit a collective type of underdetermination. Consider the following conditions.

($\forall$ Collective UD) $\mathscr{R}$ is non-empty and for each world $W_1 \in \mathscr{R}$, there is a non-isomorphic world $W_2 \in \mathscr{R}$ such that W_1 is observationally indistinguishable from W_2.

($\exists$ Collective UD) For some world $W_1 \in \mathscr{R}$ there is a non-isomorphic world $W_2 \in \mathscr{R}$ such that W_1 is observationally indistinguishable from W_2.

Clearly, any collection $\mathscr{R} \subseteq \mathscr{U}$ that satisfies ($\forall$ Collective UD) also satisfies ($\exists$ Collective UD). The implication does not go in the other direction in the sense that there are some collections $\mathscr{R} \subseteq \mathscr{U}$ that satisfy ($\exists$ Collective UD) but not ($\forall$ Collective UD). Indeed, $\mathscr{U}$ satisfies ($\exists$ Collective UD) given Proposition 2 but not ($\forall$ Collective UD) given that it contains a $\mathscr{U}$-maximal world with a God point. Such a world is observationally indistinguishable only from worlds isomorphic to it. Proposition 3 can now be recast as the following no-go result.

Corollary 3. No collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ satisfies (Heraclitus), (Leibniz), and ($\exists$ Collective UD). Thus, no collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ satisfies (Heraclitus), (Leibniz), and ($\forall$ Collective UD).

While the conditions (Heraclitus), (Leibniz), and ($\exists$ Collective UD) cannot be satisfied by any collection of worlds, any two can be for some non-empty $\mathscr{R} \subseteq \mathscr{U}$. We have already seen that if $\mathscr{R}$ is the the collection $\mathscr{H}^*$ of Heraclitus-maximal worlds, then it satisfies (Heraclitus), (Leibniz) but not ($\exists$ Collective UD). It proves instructive to work through the other two cases. Consider any Heraclitus world W_1 without a God point. Such worlds exist; see the example in Manchak and Barrett (2023). For any event $e \in W_1$, the past $P(e)$ is a proper sub-world W_1 – call it W_2. From (H1), we know W_2 is a Heraclitus world. And the two worlds are not isomorphic since, if they were, then W_2 would properly extend itself which is impossible given (H3). Because the past of any event in W_2 is contained in W_1, it follows that W_2 is observationally indistinguishable from W_1. So the collection of worlds $\mathscr{R} = \{W_1, W_2\}$ satisfies (Heraclitus) and ($\exists$ Collective UD). Of course $\mathscr{R}$ fails to satisfy (Leibniz) since W_2 is not $\mathscr{R}$-maximal.

Now consider the other case. Consider any $\mathscr{U}$-maximal world W_1 without a God point (e.g. the Minkowski world). There is a non-isomorphic $\mathscr{U}$-maximal world W_2 such that W_1 is observationally indistinguishable from W_2 (Manchak 2011, p. 418). So the collection $\mathscr{R} = \{W_1, W_2\}$ satisfies ($\exists$ Collective UD). Since any $\mathscr{U}$-maximal world must be a $\mathscr{R}$-maximal world, we see that $\mathscr{R}$ satisfies (Leibniz) as well. But since the existence of W_2 is given via the clothesline construction, we see that $\mathscr{R}$ fails to satisfy (Heraclitus). Stepping back, these two examples show a compatibility between ($\exists$ Collective UD) and either (Heraclitus) or (Leibniz). One wonders if similar compatibility results follow for ($\forall$ Collective UD) and either (Heraclitus) or (Leibniz).

We now turn to a type of underdetermination relative to the individual observer. Some worlds W are such that every individual in W, even if she lives forever, experiences observational horizons in the sense that she will never "see" some regions of W (Rindler 1956). Indeed, this is true of the de Sitter world – an important standard model of cosmology (Hawking and Ellis 1973). The de Sitter world is highly symmetric on a global, matter-averaged scale. But one can imagine worlds with causal structures similar to the de Sitter world which are also Heraclitus at a fine-grained scale. Indeed, we will construct below a collection of such worlds that, despite Corollary 3, are fully consistent with (Heraclitus), (Leibniz), and two types of individual underdetermination. Consider the following conditions on a collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ (cf. Butterfield 2014. p. 60).

($\forall$ Individual UD) $\mathscr{R}$ is non-empty and for each world $W_1 \in \mathscr{R}$, each event $e_1 \in W_1$ is such that there is a non-isomorphic world $W_2 \in \mathscr{R}$ with event $e_2 \in W_2$ such that $P(e_1)$ and $P(e_2)$ are isomorphic.

($\exists$ Individual UD) For some world $W_1 \in \mathscr{R}$, each event $e_1 \in W_1$ is such that there is a non-isomorphic world $W_2 \in \mathscr{R}$ with event $e_2 \in W_2$ such that $P(e_1)$ and $P(e_2)$ are isomorphic.

Clearly, any collection $\mathscr{R} \subseteq \mathscr{U}$ that satisfies ($\forall$ Individual UD) also satisfies ($\exists$ Individual UD). In addition, we see that ($\forall$ Collective UD) implies ($\forall$ Individual UD); and ($\exists$ Collective UD) implies ($\exists$ Individual UD). There are no other implication relations among the four conditions in the sense that one can find a counterexample in the form of some collection $\mathscr{R} \subseteq \mathscr{U}$. The collection $\mathscr{U}$ satisfies ($\exists$ Individual UD) and ($\exists$ Collective UD) but neither ($\forall$ Individual UD) nor ($\forall$ Collective UD); the collection $\mathscr{R}$ constructed in Proposition 5 below is such that it satisfies ($\forall$ Individual UD) and ($\exists$ Individual UD) but neither ($\forall$ Collective UD) nor ($\exists$ Collective UD). The entire situation is represented in the diagram below. Arrows correspond to implication relations; if two conditions in the diagram are not connected by an arrow (or series of arrows), then the corresponding implication relation does not hold.

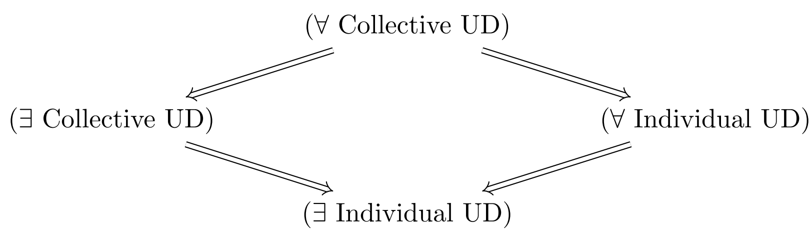

We now show that the (Heraclitus) and (Leibniz) conditions are not strong enough rule out either type of individual underdetermination. Consider the following.

Proposition 5. There is a collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ satisfying (Heraclitus), (Leibniz), and (∀ Individual UD). Thus, there is a collection of worlds $\mathscr{R} \subseteq \mathscr{U}$ satisfying (Heraclitus), (Leibniz), and (∃ Individual UD).

Proof. We will construct a collection $\mathscr{R} = \{W_1, ..., W_4\}$ with the desired properties where each of the worlds in the collection is a certain sub-world of a given Heraclitus world. Let H be the Heraclitus world in (t, x) coordinates constructed in Manchak and Barrett (2023). We note that H is conformally flat and thus has the same local causal structure as the Minkowski world. For each $i = 1, ..., 4$, let S_i be the $t > 10$ portion of the past $P(12, 2i)$. One can verify that the union $S_1 \cup ... \cup S_4$ is a sub-world of H; call this world W (see Figure 2). It is important to note that for any $i = 1, ..., 4$, any event $e \in S_i$ is such that $P(e) \subseteq S_i$. Now let $e_i = (11, 2i)$ for $i = 1, ..., 4$. We now construct four worlds as follows (see Figure 2).

$$
\begin{aligned}
W_1 &= W - \{e_1, e_2\} \\
W_2 &= W - \{e_3, e_4\} \\
W_3 &= W - \{e_2, e_3\} \\
W_4 &= W - \{e_1, e_4\}
\end{aligned}
$$

The collection of worlds $\mathscr{R} = \{W_1, ..., W_4\}$ satisfies (Heraclitus) since each world is a sub-world of a Heraclitus world. To see that (Leibniz) holds, consider any world, say W_1. The world W_1 cannot be isomorphically embedded into W_2

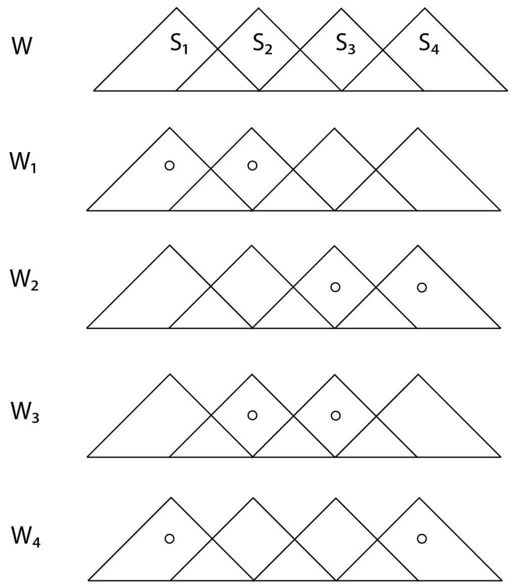

Figure 2: The world W and four of its sub-worlds $W_1, ..., W_4$.

or W_3 since the event $e_3 \in W_1$ is missing in those worlds; and W_1 cannot be isomorphically embedded into W_4 since the event $e_4 \in W_1$ is missing in that world. So W_1 is $\mathscr{R}$-maximal. The cases for worlds W_2, W_3, and W_4 are handled similarly. We also find that $\mathscr{R}$ satisfies ($\forall$ Individual UD) as well. Consider again the world W_1 and any event $e \in W_1$. Since $e \in S_i$ for some $i = 1, ..., 4$, we know $P(e) \subseteq S_i$. If e is in the S_1 region of W_1, then $P(e)$ has an isomorphic counterpart in the S_1 region of W_4; if e is in the S_2 region of W_1, then $P(e)$ has an isomorphic counterpart in the S_2 region of W_3; if e is in the S_3 region of W_1, then $P(e)$ has an isomorphic counterpart in the S_3 region of W_4; if e is in the S_4 region of W_1, then $P(e)$ has an isomorphic counterpart in the S_4 region of W_3. The cases for worlds W_2, W_3, and W_4 are handled similarly. So $\mathscr{R}$ satisfies (Heraclitus), (Leibniz), ($\forall$ Individual UD), and thus ($\exists$ Individual UD). □

A few notes about the worlds constructed in the proof above. One could have employed a more complicated construction where certain notches are removed from W instead of points so as to ensure that the worlds are extremely causally well-behaved in the sense of being "globally hyperbolic" (Hawking and Ellis 1973). Given the presence of observational horizons, such worlds have causal structures very similar to that of the de Sitter world (see Earman 1995, p. 131). Moreover, we note that an underdetermination problem would also have arisen even if we had we defined the ($\forall$ Individual UD) and ($\exists$ Individual UD) conditions using the pasts of observer world-lines instead the pasts of events (cf. Glymour 1977). Finally, given that the worlds constructed are two-dimensional, they already have an extremely well-behaved local structure in the sense that they are vacuum solutions of Einstein's equation (Fletcher et al. 2018). Stepping back, it is it hard to see what conditions, in addition to (Heraclitus) and (Leibniz), could be strong enough to secure a no-go result concerning individual underdetermination. One could perhaps object that the collection $\mathscr{R} = \{W_1, ..., W_4\}$ is artificially small in some sense. But we now use an application of Zorn's lemma to show that any collection of worlds that satisfies (Heraclitus), (Leibniz), and either (∀ Individual UD) or (∀ Individual UD) can be "extended" so as to be maximal with respect to the joint satisfaction of these conditions. This shows a type of second order satisfaction of the Leibnizian demands.

Consider the power collection $\mathcal{P}(\mathscr{U})$. Any second order property of collections of worlds gives rise to a collection $\mathcal{R} \subseteq \mathcal{P}(\mathscr{U})$ of collections of worlds with the property. We can partially order $\mathcal{P}(\mathscr{U})$ by the $\subseteq$ relation. For any collections $\mathscr{R}, \mathscr{R}' \in \mathcal{P}(\mathscr{U})$, we say that $\mathscr{R}'$ is a (not necessarily proper) **extension** of $\mathscr{R}$ if $\mathscr{R} \subseteq \mathscr{R}'$. For any second order property $\mathcal{R} \subseteq \mathcal{P}(\mathscr{U})$, we say that a collection $\mathscr{R} \in \mathcal{R}$ is **$\mathcal{R}$-maximal** if it has no proper extensions in $\mathcal{R}$. One has a general question: which second order properties $\mathcal{R} \subseteq \mathcal{P}(\mathscr{U})$ are such that any collection in $\mathcal{R}$ has an $\mathcal{R}$-maximal extension? Some cases are trivial to settle. Consider the collection $\mathcal{H} \subset \mathcal{P}(\mathscr{U})$ of all collections of worlds satisfying the (Heraclitus) condition. Any collection of worlds in $\mathcal{H}$ has an $\mathcal{H}$-maximal extension: the collection $\mathscr{H}$ of all Heraclitus worlds. The four underdetermination conditions are handled similarly. Consider (∀ Collective UD) for example. Let $\mathcal{C} \subset \mathcal{P}(\mathscr{U})$ be the collection of all collections of worlds satisfying the (∀ Collective UD) condition. One can verify that the union $\bigcup \mathcal{C}$ satisfies (∀ Collective UD) as well and therefore counts as a $\mathcal{C}$-maximal extension for any collection of worlds in $\mathcal{C}$.

The (Leibniz) case is a bit more interesting. Consider the collection $\mathcal{L} \subset \mathcal{P}(\mathscr{U})$ of all collections of worlds satisfying the (Leibniz) condition. The union $\bigcup \mathcal{L}$ is too big to satisfy (Leibniz). To see this, consider any Heraclitus world W_1 and any of its proper sub-worlds W_2. As we have seen, the collections $\{W_1\}$ and $\{W_2\}$ satisfy (Leibniz). So $W_1, W_2 \in \bigcup \mathcal{L}$ which implies that $\bigcup \mathcal{L}$ does not satisfy (Leibniz). But now consider any collection $\mathcal{T} \subset \mathcal{L}$ totally ordered by the $\subseteq$ relation. The union $\bigcup \mathcal{T}$ is an upper bound for $\mathcal{T}$. It also satisfies (Leibniz). If it didn't, there would be worlds $W_1, W_2 \in \bigcup \mathcal{T}$ such that one is a proper extension of the other. But this cannot be since it implies that W_1 and W_2 can be found in some collection in $\mathcal{T}$ and all such collections satisfy (Leibniz). From Zorn's lemma, it now follows that any collection in $\mathcal{L}$ has an $\mathcal{L}$-maximal extension. A similar argument shows the following. (It is an open question whether Zorn's lemma is needed here.)

Proposition 6. Let $\mathcal{R} \subset \mathcal{P}(\mathscr{U})$ be the collection of all collections of worlds satisfying (Heraclitus), (Leibniz), and either of the following: (∀ Individual UD) or (∃ Individual UD). Then any collection $\mathscr{R} \in \mathcal{R}$ has an $\mathcal{R}$-maximal extension.

7 Conclusion

Here, we have introduced the notion of Heraclitus-maximal worlds within the context of general relativity. Heraclitus worlds are characterized by an asymmetry property: no distinct events in a world have the same structure. Leibnizian metaphysics is often deployed to consider worlds which are maximal with respect to their properties. A Heraclitus-maximal world is a world that is "as large as it can be" with respect to the Heraclitus property. We have investigated some basic properties of Heraclitus-maximal worlds. In particular, we have shown that they exist and that any Heraclitus world can be extended to a Heraclitus-maximal world. Finally, we have considered the following question: can one know which world one inhabits? For collections of worlds satisfying the (Heraclitus) and (Leibniz) conditions, there are senses of underdetermination in which the answer must be yes (universal and existential collective types) and others in which the answer can be no (universal and existential individual types).

We close with a question. As we have seen, the results here depend very little on the structure of general relativity. Consider any spacetime theory given by a collection of worlds where each world is represented by a manifold with some geometric structure defined on it. For example, consider a classical spacetime theory of this type (Malament 2012). In the natural way, one can define subworlds and isomorphisms between worlds in this theory. With these notions, one can go on to define Heraclitus-maximal worlds and explore their basic properties. One can also articulate the analogs of the conditions (Heraclitus) and (Leibniz). The relation of observational indistinguishability among worlds and the various types of underdetermination conditions can also be introduced and explored. One wonders: do the results considered here carry over to the new context?

References

[1] Barrett, T. (2018), "What Do Symmetries Tell Us About Structure?" *Philosophy of Science*, 85: 617-39.

[2] Belot, G. (2017), "Fifty Million Elvis Fans Can't be Wrong," *Noûs*, 52: 946-981.

[3] Belot, G. (2023), *Accelerating Expansion: Philosophy and Physics with a Positive Cosmological Constant*, Oxford: Oxford University Press.

[4] Bradley, C. and Weatherall, J. (2020), "On Representational Redundancy, Surplus Structure, and the Hole Argument," *Foundations of Physics*, 50: 270-293.

[5] Butterfield, J. (2014), "On Under-Determination in Cosmology," *Studies in History and Philosophy of Modern Physics*, 46: 57-69.

[6] Cinti, E., and Fano, V. (2021) "Careful with Those Scissors, Eugene! Against the Observational Indistinguishability of Spacetimes," *Studies in History and Philosophy of Modern Physics*, 89: 103-113.

[7] Clarke, C. (1993), *The Analysis of Space-Time Singularities*, Cambridge: Cambridge University Press.

[8] Dasgupta, S. (2016), "Symmetry as an Epistemic Notion (Twice Over)" *The British Journal for the Philosophy of Science*, 67: 837-878.

[9] Earman, J. and Norton, J. D. (1987), "What Price Spacetime Substantivalism? The Hole Story," *The British Journal for the Philosophy of Science*, 38: 515-525.

[10] Earman, J. (1989), *World Enough and Space-Time*, Boston: MIT Press.

[11] Earman, J. (1995), *Bangs, Crunches, Whimpers, and Shrieks: Singularities and Acausalities in Relativistic Spacetimes*, Oxford: Oxford University Press.

[12] Fletcher, S. (2020) "On Representational Capacities, with an Application to General Relativity," *Foundations of Physics*, 50: 228-249.

[13] Fletcher, S., Manchak, J., Schneider, M., and Weatherall, J. (2018), "Would Two Dimensions Be World Enough for Spacetime?" *Studies in History and Philosophy of Modern Physics*, 63: 100-113.

[14] Geroch, R. (1970), "Singularities," in M. Carmeli, S. Fickler, and L. Witten (eds.), *Relativity*, New York: Plenum Press, 259-291.

[15] Glymour, C. (1972), "Topology, Cosmology, and Convention," *Synthese*, 24: 195-218.

[16] Glymour, C. (1977), "Indistinguishable Space-Times and the Fundamental Group," in J. Earman, C. Glymour, and J. Stachel, eds., *Foundations of Space-Time Theories, Minnesota Studies in the Philosophy of Science, vol 8*, Minneapolis: University of Minnesota Press, 50-60.

[17] Hawking, S., and Ellis, G. (1973), *The Large Scale Structure of Space-Time*, Cambridge: Cambridge University Press.

[18] Hebda, J. (2010), "Heterogeneous Riemannian Manifolds," *International Journal of Mathematics and Mathematical Sciences*, 2010: 187232.

[19] Krasnikov, S. (2018), *Back-in-Time and Faster-than-Light Travel in General Relativity*, Cham: Springer Nature.

[20] Low, R. (2012), "Time Machines, Maximal Extensions and Zorn's Lemma," *Classical and Quantum Gravity*, 29: 097001.

[21] Malament, D. (1977), "Observationally Indistinguishable Space-Times," in J. Earman, C. Glymour, & J. Stachel (eds.), *Foundations of Space-Time Theories, Minnesota Studies in the Philosophy of Science Vol. VIII*, Minneapolis: University of Minnesota Press, 61-80.

[22] Malament, D. (2012), *Topics in the Foundations of General Relativity and Newtonian Gravitation Theory*, Chicago: University of Chicago Press.

[23] Manchak, J. (2009), "Can We Know the Global Structure of Spacetime?" *Studies in History and Philosophy of Modern Physics*, 40: 53-56.

[24] Manchak, J. (2011), "What is a Physically Reasonable Spacetime?" *Philosophy of Science*, 78: 410-420.

[25] Manchak, J. (2014), "On Space-Time Singularities, Holes, and Extensions" *Philosophy of Science*, 81: 1066-1076

[26] Manchak, J. (2016), "Is the Universe As Large as It Can Be?" *Erkenntnis*, 81:1341-1344.

[27] Manchak, J. (2021), "General Relativity As a Collection of Collections of Models," in J. Madarász and G. Székely eds., *Hajnal Andréka and Istvań Németi on Unity of Science: from Computing to Relativity Theory Through Algebraic Logic*, Springer Nature, 409-425.

[28] Manchak, J. (2023), "On the (In?)Stability of Spacetime Inextendibility," *Philosophy of Science*, 90: 1331-1341.

[29] Manchak, J. and Barrett, T. (2023), "A Hierarchy of Spacetime Symmetries: Holes to Heraclitus," forthcoming in *The British Journal for the Philosophy of Science*.

[30] North, J. (2021), *Physics, Structure, and Reality*, Oxford: Oxford University Press.

[31] Norton J. (2011), "Observationally Indistinguishable Spacetimes: A Challenge for Any Inductivist". In G. Morgan, ed., *Philosophy of Science Matters: The Philosophy of Peter Achinstein*, Oxford: Oxford University Press, 164-76.

[32] O'Neill, B. (1983), *Semi-Riemannian Geometry with Applications to Relativity*, London: Academic Press.

[33] Penrose, R. (1969), "Gravitational Collapse: The Role of General Relativity," *Revisita del Novo Cimento, Serie I*, 1: 252-276.

[34] Rindler, W. (1956), "Visual Horizons in World-Models," *Monthly Notices of the Royal Astronomical Society*, 116: 662-677.

[35] Roberts, B. (2020), "Regarding 'Leibniz Equivalence'," *Foundations of Physics*, 50: 250-269.

[36] Sbierski, J. (2016), "On the Existence of a Maximal Cauchy Development for the Einstein Equations: a Dezornification," *Annales Henri Poincaré*, 17: 301-329.

[37] Stein, H. (1967), "Newtonian Space-Time," *Texas Quarterly*, 10: 174-200.

[38] Sunada, T. (1985), "Riemannian Coverings and Isospectral Manifolds," *Annals of Mathematics*, 121: 169-186.

[39] Weatherall, J. (2018), "Regarding the 'Hole Argument"', *The British Journal for the Philosophy of Science*, 69: 329-350.

[40] Weyl, H. (1952), *Symmetry*, Princeton: Princeton University Press.

[41] Wong, W. (2013), "A Comment On the Construction of the Maximal Globally Hyperbolic Cauchy Development," *Journal of Mathematical Physics*, 54: 113511.