Belief in Robust Temporal Passage (Probably) Does Not Explain Future-Bias

Abstract

Empirical work has lately confirmed what many philosophers have taken to be true: people are 'biased toward the future'. All else being equal, we usually prefer to have positive experiences in the future, and negative experiences in the past. According to one hypothesis, *the temporal metaphysics hypothesis,* future-bias is explained either by our (tacit) beliefs about temporal metaphysics—the *temporal belief hypothesis*—or alternatively by our temporal phenomenology—the *temporal phenomenology hypothesis*. We empirically investigate a particular version of the temporal belief hypothesis according to which future-bias is explained by the belief that *time robustly passes*. Our results do not match the apparent predictions of this hypothesis, and so provide evidence against it. But we also find that people give more future-biased responses when asked to *simulate* a belief in robust passage. We take this to suggest that the phenomenology that attends simulation of that belief may be partially responsible for future-bias, and we examine the implications of these results for debates about the rationality of future-bias.

1. Introduction

Philosophers have long supposed, going back at least to Hume (1739), that humans exhibit a *bias toward the future*: At least in some contexts, we prefer that positive events (events that we prefer to occur) are located in the future rather than the past, and that negative events (events that we prefer not to occur) are located in the past rather than the future. At a minimum, when it comes to our own pleasant and unpleasant experiences, we usually prefer that pleasant experiences be located in the future and unpleasant experiences in the past.

Call events involving pleasant or painful experiences *hedonic events,* and all other events *non-hedonic events*. So, for instance, experiencing being poked by a hot poker is a (negative) hedonic event, whilst unknowingly being insulted by a friend behind your back

is a (negative) non-hedonic event. In these terms, philosophers have generally supposed that we are future-biased with respect to hedonic events, but not non-hedonic events.[1]

We will say that someone is positively hedonically future-biased if, all else being equal, they tend to prefer positive hedonic events in their future rather than their past, and that they are negatively hedonically future-biased if they tend to prefer negative hedonic events to be in their past rather than their future. Recent work supports the contention that people are both positively and negatively hedonically future-biased (Caruso, Gilbert and Wilson 2008; Greene, Latham, Miller and Norton 2020) in a first-person setting. That is, people show future-bias when they are asked to consider their preferences over their own hedonic experiences, rather than over the experiences of some third party. By contrast, evidence about patterns of future-bias in third-person conditions—where people are asked about their preferences over the temporal location of hedonic events for some third-party—is equivocal. Caruso et al. (2008), using a study design in which they ask participants how much compensation ought to be paid to themselves, or others, for work either undertaken in the past or in the future, found that people awarded more compensation for future work only in the first-person setting. These results are usually interpreted as showing that people exhibit time-neutral preferences in the third-person but not first-person condition. By contrast, using a somewhat different methodology that explicitly asked participants about their preferences, Greene et al. (2020) found that most people are future-biased about hedonic events in both first-person and third-person conditions. Given the disparity in these findings, in this paper we focus solely on first-person hedonic conditions, in which there have been robust findings of future-bias.

A question that remains largely empirically unexplored, however, is the psychological explanation for future-bias. Philosophers have suggested at least two candidate explanations. The first is what we call the *temporal metaphysics hypothesis*. Very broadly, according to this hypothesis, our beliefs (likely implicit) about the temporal structure of our world (the *temporal belief hypothesis*) and/or our temporal phenomenology (the *temporal phenomenology hypothesis*) explain our future-biased preferences. The second is what we call the *practical irrelevance hypothesis*. According to this hypothesis, we are future-biased because there is nothing we can do to affect the past, which means that past events

[1]See for instance Prior (1959), Hare (2007, 2008), and Heathwood (2008).

are not relevant to our choices (and therefore are not objects of *practical* concern) in the way that potential future events are.[2]

These two explanatory hypotheses have potential implications for the question of whether our future-biased pattern of preferences is *rational*. There has recently been considerable debate over the rationality of future-bias (Hare 2007, 2013; Brink 2011; Dougherty 2011, 2015; Greene & Sullivan 2015; Sullivan 2018; Dorsey 2018). Some of this discussion appeals to our (purported) *pattern* of future-biased preferences. In particular, it has been argued that because we have future-biased preferences in first-person hedonic conditions and not in third-person conditions or non-hedonic conditions, our future-biased preferences are being driven by irrational considerations.[3] Another way to evaluate the rationality of future-bias, however, is to focus on the underlying psychological factors that explain it. For instance, one might think that the causal inaccessibility of the past is a good enough reason to care less about one's past experiences; but this only rationalises our *actual* future-biased preferences if the practical irrelevance hypothesis is correct, i.e., if the belief that the past is causally inaccessible is in fact what explains our future-bias. On the other hand, if some version of the temporal metaphysics hypothesis is true, then the rationality of (actual) future-bias seems to turn on whether our beliefs about the metaphysics of time are justified, or our experience of time veridical, and whether those beliefs/experiences in fact provide a reason to care more about the future than the past.

This paper empirically investigates one version of the temporal metaphysics hypothesis (and, more specifically, the temporal belief hypothesis), which we will call the *passage belief hypothesis*. According to this hypothesis, our future-biased preferences are (at least partly) explained by our belief *that time robustly passes*. We aim to test this hypothesis. In §2 we give a more detailed survey of the relevant literature on future-bias, and present our hypotheses. In §3 we describe our methodology and results. In §4 we consider the implications of those results for our understanding of why we have future-biased preferences and, in turn, for arguments concerning the rationality of those preferences.

[2] Defenders of something like this view include Kauppinen (2018, 1984, p. 186) and Horwich (1987, pp. 194-196). It is developed more fully and explicitly by Maclaurin & Dyke (2002) and Suhler & Callender (2012).

[3] Brink argues that non-hedonic time-neutrality would show that future-bias is limited in scope and therefore suspect as a component of a theory of rationality. Dougherty (2015: p. 3, fn. 4) interprets this assumed nonexistence of non-hedonic future-bias to be evidence of the arbitrariness of any future-biased preferences, since it suggests that future-biased preferences are not formed by rational processes. However, this purported pattern of preferences was not, in fact, vindicated by Greene et al. (2020) who found no asymmetry between future-bias with regard to hedonic as opposed to non-hedonic events.

2. The Literature and Our Hypotheses

While there has been very little experimental work on the explanation for future-bias, recent work by Latham, Miller, Norton and Tarsney (2020) appears to provide some empirical support for the practical irrelevance hypothesis: future-bias is mitigated when people are asked to consider hypothetical choice situations where they can causally influence past events, suggesting that it is at least in part the practical irrelevance of the past that is responsible for future-bias. Importantly, though, their study found significant residual future-bias even when the past is causally accessible. This suggests that the practical irrelevance of the past provides only a partial explanation for future-bias, and hence that some version of the temporal metaphysics hypothesis might play a complementary explanatory role.

As noted, there are at least two versions of the temporal metaphysics hypothesis: the temporal belief hypothesis (according to which future bias is at least partly explained by our beliefs about temporal metaphysics) and the temporal phenomenology hypothesis (according to which future bias is at least partly explained by our temporal phenomenology). Both of these more specific hypotheses themselves have several versions.

An especially natural version of the temporal phenomenology hypothesis holds that it is the phenomenology as of *robust temporal passage* that (at least partly) explains our future-biased preferences. By 'robust temporal passage', we mean the kind of passage posited by A-theories of time. A-theories of time hold that there is an objective, observer-independent fact about which moment (or set of events) is present, and which moment (or events) this is changes. Robust temporal passage is just this change in which events are objectively present.[4]

B-theories of time, by contrast, hold that moments in time only stand in *earlier than* and *later than* relations to one another, and are not 'past', 'present', or 'future' except from the perspective of particular observers located at particular times. Since there is no objective present, there cannot be robust passage.[5] A 'phenomenology as of temporal passage', then, is a phenomenology whose content represents that the world contains

[4] Sometimes this change consists in the movement of a property of presentness across existing events, which then change from being future, to being present, (when they have the property) to being past (as in a moving spotlight model). Sometimes this change consists simply in the change of a single three-dimensional object with respect to which objects or events exist (as in presentism) and sometimes it consists in the accretion of new moments of time or events, where these new moments/events are present when they come into existence, and then become past as new moments/events come into existence.

[5] The terminology of 'A-theory' and 'B-theory' is originally due to McTaggart (1908). For an overview of the debates between A-theorists and B-theorists, see Zimmerman (2005).

robust passage, whether or not the world is in fact the way it is represented to be (hence the *as of)*. This phenomenology might involve, among other things, the feeling that future events are moving towards us (or, alternatively, that we are moving towards them) and that past events are moving away from us (or, alternatively, that we are moving away from them). The idea, very roughly, is that this phenomenology contributes to future-bias because, if an event is positively valenced, we will prefer to feel as though we are getting closer to that event, and hence prefer it to be in the future, whereas if the event is negatively valenced, we will prefer to feel as though we are getting further away from it, and hence prefer it to be in the past. Call this the *passage phenomenology hypothesis*. This is just one of various hypotheses according to which some aspect of our temporal phenomenology contributes to future-bias. For present purposes, however, we set aside all versions of the temporal phenomenology hypothesis and focus entirely on the temporal belief hypothesis (though we return to the temporal phenomenology hypotheses in section 4).

As with the temporal phenomenology hypothesis, there are many possible versions of the temporal belief hypothesis, corresponding to different beliefs about time that might be thought to play a role in explaining future-bias. For instance, there is what we might call the *temporal ontology belief hypothesis*. According to this hypothesis, our (tacit) beliefs about temporal ontology play a role in explaining our future-bias. On the one hand, if you believe that past and future events have the same ontological status (e.g., all being part of the same 'frozen block' of events), then there is a *prima facie* case for neutrality between your past and future selves. On the other hand, if you believe that there is some asymmetry in ontological status between past and future events (or between your past and future selves), this might give you reason to have greater concern for the future. Perhaps, for instance, future events are uniquely *open* or *indeterminate,* or have some feature like *potential* or *incipient* actuality that past events lack. Even if both past and future events lack the sort of reality enjoyed by present events, it might seem rational to care more about the future if it is 'in the nature of' future events (but not past events) to *become* present. If most people believe in such an ontological asymmetry, this could explain why most people are future-biased.[6]

There are reasons to be sceptical of this hypothesis, though. In particular, it's not obvious that folk beliefs about time involve any past-future ontological asymmetry of the sort that might explain/justify future-bias. If anything, the commonly held 'growing block' view, according to which the past is real but the future isn't, seems intuitively to support

[6] See Latham, Miller and Norton (2019) for a discussion of which views about temporal ontology are most prevalent amongst the folk.

past-bias, i.e., caring more about past experiences than future experiences. And if the past and future are both simply non-existent, this does not in itself give us any asymmetry to work with at all. Meanwhile, the belief that the future is *potentially actual* or *will become actual* is arguably not an ontological belief at all but just a way of describing a belief in robust passage.[7]

The hypothesis we aim to investigate in the present study, therefore, is a different version of the temporal belief hypothesis—namely, the passage belief hypothesis, according to which a (tacit) belief that time robustly passes contributes to future-bias. As we have already said, to believe in robust temporal passage is to believe that there is an objective fact about which events are present, and that which events those are changes, either as new events come into existence, or as future events become present (and then past), or as present events cease to exist and new ones come into existence. By contrast, if one believes that events simply stand in earlier/later relations without any times or events being picked out as objectively past, present, or future, then one believes that there is no such thing as robust passage.[8]

The passage belief hypothesis has been defended by a number of authors, but seems to have been inspired by Prior (1959), who argues that verbal expressions of time-asymmetric attitudes like relief require an irreducibly tensed semantics, which in turn implies (the speaker's belief in) a metaphysics of irreducibly tensed facts or properties.[9] Others have argued for a link between future-bias and belief in robust temporal passage more directly, without going via semantics—for instance, Schlesinger (1976) and Craig (1999).

It is easy to see why some philosophers have found this hypothesis plausible. If one believes in robust temporal passage, it seems natural to say that there is some good sense in which either future events are 'coming towards us' (as we remain located at the objective present) and then recede ever further away from us into the past, or we are 'moving towards the future and away from the past', carried along with the moving present,

[7] Thanks to an anonymous referee for helpfully highlighting these limitations of the temporal ontology belief hypothesis.

[8] Of course, believing that time is static is not the same as believing that there are no temporal asymmetries. Those who defend static views of time hold either that time itself is asymmetric, or at the very least, that there are asymmetries in time (such as the asymmetry of causation, thermodynamic asymmetries, epistemic asymmetries, and so on). So it could be that belief in *these* asymmetries plays a role in explaining future-bias. If so, we would expect *these* factors to still be at play amongst people who believe that time is static.

[9] Prior's argument has been widely criticized (most famously by Mellor (1981, 1983) and MacBeath (1983)), but remains influential. For a recent defence of Prior's conclusion against the Mellor/MacBeath reply, see Pearson (2018).

as events themselves remain fixed in time. The sorts of locutions that would capture robust passage are, indeed, ones that we find frequently in ordinary language (at least, in English). It has been argued that people's understanding of time is largely structured around the use of spatial metaphors (Lakoff & Johnson 1980; Boroditsky 2000). Common expressions in everyday language tend to reflect, or at least suggest, a view of time as dynamic, with past events 'behind' the self and growing more distant (e.g., 'The worst is behind us'), and future events 'ahead of' the self and growing closer (e.g., 'She has a bright future ahead of her') (Boroditsky 2000; Clark 1973). And, rightly or wrongly, it seems natural to care more about things that are 'moving towards us' than things that 'moving away from us'—at a minimum, the former are better at commanding our attention.

The passage belief hypothesis and the passage phenomenology hypothesis may, of course, be connected. It might be that we tacitly believe that time robustly passes *because* we have a phenomenology as of time robustly passing.[10] In turn, if this belief is what explains our future-bias, then it seems equally correct to say that the corresponding phenomenology explains our future-bias. Then both versions of the temporal metaphysics hypothesis would be vindicated. Equally, though, either passage phenomenology or passage beliefs could give rise to future-bias without any involvement from the other (most clearly, in the case that we have passage beliefs without any corresponding phenomenology or vice versa). In this paper we do not try to assess the potential connections between belief-based and phenomenology-based hypotheses: we focus only on the question of whether the passage belief hypothesis has any empirical support.

In order to probe this question, we ran a study in which we divided participants between two conditions. In the first condition, participants see a vignette in which they are asked to imagine that they are an astronaut (in space) who has discovered that time robustly passes. This is the *passage condition.* In the second condition, participants see a vignette in which they are asked to imagine that they are an astronaut (in space) who has discovered that time does not robustly pass. This is the *no-passage condition.* In each condition participants are told that the spaceship's food dispenser normally produces bland meals, but that once during the voyage it produces the astronaut's favourite meal (the *positive hedonic condition*) or their most disliked meal (the *negative hedonic condition*). Participants are then asked questions about their preferences regarding the temporal location of their favourite/most-disliked meal.

[10] Some defenders of robust passage take it that the content of our phenomenology gives us reason to believe that time passes. See Smith (1994) and Craig (1999).

If believing that time robustly passes is a partial explanation for future-biased preferences, then we would expect that more participants in the passage condition will exhibit future-biased preferences than will participants in the no-passage condition. This was our first hypothesis (hypotheses were preregistered at https://osf.io/r2g5e/[11]).

> H1: More participants in the passage condition will exhibit future-biased preferences than will participants in the no-passage condition.

Second, we would expect to find that amongst those who report future-biased preferences across the passage and no-passage conditions, that the strength of those future-biased preferences is stronger in the passage condition than in the no-passage condition. This is our second hypothesis:

> H2: Future-biased preferences will be stronger in the passage condition than in the no-passage condition.

Importantly, our vignettes do not directly impact agents' actual beliefs about robust passage. Rather, they ask participants to imagine, or simulate, that time either does or does not robustly pass. So, even if it's the case that people tend to be more future-biased when they simulate having the belief that time robustly passes, it doesn't follow that people's actual beliefs about robust passage are an explanation for their future-biased preferences. (For instance, it might be that no one actually believes in robust passage.) Since we want to know what explains people's *actual* future-biased preferences, we want to know not only whether *simulating* the belief that time robustly passes is associated with greater future-bias, but whether *actually believing* that time robustly passes is associated with greater future-bias.

So, in order to more directly assess the hypothesis we're interested in, we asked participants in the passage and no-passage conditions how likely they think it is that our universe matches the description in the vignette (i.e., contains/doesn't contain robust passage). We take this to be a way to probe their (possibly tacit) beliefs about whether time does in fact robustly pass or not. We need to be careful here, though: a participant who does *not* believe that the universe is as the passage condition describes need not believe that the universe *is* as the no-passage condition describes (and likewise, a participant who

[11] [This link will only become accessible once the paper is de-anonymised, since the information at the link effectively identifies the author(s) of the paper.]

does not believe that the universe is as the no-passage condition describes need not believe that the universe is as the passage condition describes). Given this, to describe participants' reported temporal beliefs, we will say that a participant

- *believes that time robustly passes* if they are in the passage condition, and they respond that they believe that our world is as described by that condition.
- *does not believe that time robustly passes* if they are in the passage condition, and they respond that they believe that our world is *not* as described by that condition.
- *believes that time is static* if they are in the no-passage condition, and they respond that they believe that our world is as described by that condition.
- *does not believe that time is static* if they are in the no-passage condition, and they respond that they believe that our world is not as described by that condition.

The passage belief hypothesis seems to predict that participants who either believe that time robustly passes or do not believe that time is static will exhibit more future-bias than those who do not believe that time robustly passes or believe that time is static. But it also seems natural to expect that the association between future-bias and temporal passage beliefs will be stronger in participants who believe that time robustly passes than in those who merely do not believe that time is static. That is, we should predict that participants with either of these two beliefs show higher levels of future-bias than participants with the opposite beliefs, but that the association will be stronger amongst those who believe that time robustly passes than amongst those who do not believe that time is static. That is because there are ways that time could fail to be static, which are not also ways in which time robustly passes. Hence we predicted:

> H3: The association between condition and future-bias will be stronger among people who believe that our world is as is described by the condition, than among people who believe that the world is not as described by the condition. More people who believe that time robustly passes will have future-biased preferences than those who do not. Fewer people who believe that time is static will have future-biased preferences than those who do not.

Finally, we predicted that amongst participants who have future-biased preferences, those who believe that time robustly passes will have stronger preferences than those who do not. More carefully:

H4: Among participants in a given condition who have future-biased preferences, there will be an association between believing our world is like the condition described and preference strength. People who believe that time robustly passes will have stronger future-biased preferences than those who do not. People who believe that time is static will have weaker future-biased preferences than those who do not.

3. Experimental Design and Results

3.1 Method

3.1.1 Participants

694 people participated in the study. Participants were U.S. residents, recruited and tested online using Amazon Mechanical Turk, and compensated $0.50 for approximately 5 minutes of their time. Given recent worries about the quality of data collected through MTurk, concerning both the quality of human responders and the presence of bots, we adopted a number of quality control measures.[12]

First, we used only those MTurk participants who have a HIT (task) approval rate of at least 95% and who have had their HITs (tasks) approved at least 1000 times. That means that all our participants had already successfully completed at least 1000 other studies, and received at least a 95% approval rating on these tasks, a standard that can be expected to eliminate most bots.

Second, our study included both task instructions and attentional checks that doubled as comprehension checks. We excluded participants who failed either to follow instructions or to correctly answer an attentional check/comprehension question. In total, 433 participants were excluded for either failing to answer all the questions, or failing either the attentional check or comprehension question. The remaining sample was composed of 261 participants (aged 20-74; 119 female; mean age 39.87 (SD = 12.48)).[13] Ethics approval was obtained from the [blanked] Human Research Ethics Committee. Informed consent was obtained from all participants prior to testing. The survey was conducted online using Qualtrics.

[12] See Ahler, Roush & Soud (ms) for a discussion of some of the problems associated with collecting data using MTurk and the prevalence thereof.

[13] If participants were simply responding at chance to both the attentional check question and comprehension question, then we would expect a remaining sample of 44 participants (bracketing those participants that failed to follow task instructions).

3.1.2 Materials and Procedure

This study examined the relationship between beliefs about robust temporal passage and future-bias. Participants were first randomly assigned to read one of four vignettes:

Passage Condition Positive/Negative Hedonic Event

The date is June 1, 2090. You are an astronaut on a 10-year voyage between planets. Your background is in physics, and your mission is to investigate certain aspects of time. You have already discovered that there are objective facts about which moment is truly present, and which moments are past, and which are future. It is *not* the case that events just have different locations in time, with none of them singled out as truly present. Instead, which moment is *really* present changes, so that events begin by being future, gradually come ever closer to the present, become present, and then recede into the past.

You are 5 years into the voyage. The ship's food dispenser normally produces bland meals containing only essential nutrients. However, it is programmed to dispense your **[favourite]/[most disliked]** meal — which you really **[like]/[dislike]** — during one day of the voyage. One morning, you awake from a dream concerning your **[favourite]/[most disliked]** meal and for a moment you cannot remember whether you have received it yet.

You realise that if your meal was dispensed yesterday, then your eating of that meal is now in the objective past, and is receding ever further into the past, away from the present moment. If your meal is to be dispensed tomorrow, then your eating of that meal will soon be in the objective present, and is growing ever closer to you as future events (such as the arrival of your meal) move from being located in the future, to being located in the present, to, eventually, being located in the past.

No Passage Condition, Positive/Negative Hedonic Events

The date is June 1, 2090. You are an astronaut on a 10-year voyage between planets. Your background is in physics, and your mission is to investigate certain aspects of time. You have already discovered there are no objective facts about which moment is truly present, and which moments are past and which are future. It is *not* the case that which moment is *really* present changes, so that events begin by being future, gradually come ever closer to the present, become present, and then recede into the past. Instead, all events are simply laid out in space-time and none of them is singled out as being truly present.

You are 5 years into the voyage. The ship's food dispenser normally produces bland meals containing only essential nutrients. However, it is programmed to dispense your **[favourite]/[most disliked]** meal — which you really **[like]/[dislike]** — during one day of the voyage. One morning, you awake from a dream concerning your [**favourite]/[most disliked]** meal and for a moment you cannot remember whether you have received it yet.

> You realise that if your meal was dispensed yesterday, then your eating of that meal is located one day away from your current location in space-time. Since your current location, and the location of your meal yesterday, never change their location in space-time, it will always be the case that the eating of yesterday's meal is one day away from your current location in time. If your meal is to be dispensed tomorrow, then your eating of that meal is also located one day away from your current location in time (i.e. one day away from June 1, 2090 in time). Since your current location, and the location of your meal tomorrow never change their location in time, it will always be the case that the eating of tomorrow's meal is one day away from your current location—i.e.one day away from June 1, 2090, in time.

After reading the vignette, participants were asked to indicate their preference using the following statements:

(a) 'I would prefer to learn that my favourite/most disliked meal was dispensed yesterday, and will not be dispensed tomorrow.'

(b) 'I would prefer to learn that my favourite/most disliked meal will be dispensed tomorrow, and was not dispensed yesterday.'

(c) 'I have no preference regarding what I learn about when my favourite/most disliked meal is dispensed.'

We presented participants with both (a) and (b) in order to control for question effects—effects due to being asked to agree that one would prefer the event to be in the near past versus being asked to agree that one would prefer the event to be in the near future. And (c) served to distinguish people who strongly disagreed that they had the relevant preference because they simply lacked a preference on the matter from those who disagreed because they had an opposing preference. The statements were presented in random order. Participants were also asked to indicate the strength of their preferences on a Likert scale from 1 (very weak) to 7 (very strong).

Next, we measured participants' (possibly tacit) beliefs regarding *actual* robust passage. To do this, we reminded participants which discovery had been made about time according to the vignette they read, and then asked them: 'How likely do you think it is

that our universe is like this?' Participants responded on a Likert scale from 1 (completely unlikely) to 7 (completely likely).

We also required participants to respond to comprehension and attention check questions. We asked participants: 'In the preceding vignette it was **June 2090** on the spacecraft. This means that events that take place in **August 2090** are...', to which they could answer (a) 'objectively in the future', (b) 'moving closer to the present', (c) 'both objectively in the future and moving closer to the present', and (d) 'neither objectively in the future nor moving closer to the present'. Participants in the passage condition were excluded unless they answered (c) and participants in the no-passage condition were excluded unless they answered (d).[14] Finally, participants were asked: 'In this vignette, you were asked to imagine that you were...', to which they could answer: (a) 'A Spaceship', (b) 'An Astronaut', (c) 'A Food Dispenser', or (d) 'A Dog'. Participants who did not choose (b) were excluded.

1.2. Results and Analyses

Before we summarise the main findings, let's introduce some terminology. To describe the time biases reported in our survey, we will say that a participant

- has a *positive future-biased preference* if they prefer their favourite meal to be located in the future rather than the past.
- has a *positive past-biased preference* if they prefer their favourite meal to be located in the past rather than the future.
- is *positively time-neutral* if they have no preference between these options.
- has a *negative future-biased preference* if they prefer their most disliked meal to be located in the past rather than the future.
- has a *negative past-biased preference* if they prefer their most disliked meal to be located in the future rather than the past.

[14] We found no evidence of a difference between the condition a participant was assigned to and passing the comprehension question.

- is *negatively time-neutral* if they have no preference between these options.

Before reporting the statistics, we summarise our main findings with respect to each of our hypotheses. Our hypotheses were, first, that more people would report a future-biased preference (positive or negative) in the passage condition than in the no-passage condition; second, that future-biased preferences would be stronger in the passage condition than in the no-passage condition; third, that the association between condition and future-bias will be stronger among people who believe that our world is like the condition described than among people who do not.; and fourth, that people who believe that time robustly passes or do not believe that time is static would report stronger future-biased preferences than people who do not believe that time robustly passes or believe that time is static.

Only the first two hypotheses were vindicated. More people reported being future-biased in the passage condition than in the no-passage condition, and people who reported future-bias in the passage condition had stronger future-biased preferences than those who reported future-bias in the no-passage condition. The second two hypotheses were not vindicated. We found no evidence of an effect of people's beliefs on the association between condition and future-biased preference, nor on the strength of people's future-biased preference.

Table 1 below summarises the descriptive data of participants' reported preferences across all conditions. The 'FB' column represents the number of participants who report a positive or negative future-biased preference. The 'non-FB' column represents the number of participants who report either (i) a positive or negative past-biased preference or (ii) a positive or negative time-neutral preference. We have combined these numbers due to the low numbers of participants who report having a preference that is not future-biased. Beneath the numbers (and proportions) of people who report a FB preference or a non-FB preference we report the means (and standard deviations) of people's strength of preference across all conditions. We also include the results of one-way chi-square tests, which tests, for each condition whether most people responded as future-biased. The results of these tests show that the majority of people in each condition report either a positive future-biased preference or negative future-biased preference.[15]

[15] Non-FB numbers are made up as follows: (i) positive event passage condition – 6 positive past-biased preferences and 6 positive time-neutral preferences; (ii) positive event no passage condition – 10 positive past-biased preferences and 14 positive time-neutral preferences; (iii) negative event passage condition – 9

Table 1. Descriptive data from all conditions of participants preferences.

Condition	FB	Non-FB	χ^2	*p*-value
Positive Event				
Passage (*n* = 65)	53 (81.5%)	12 (18.5%)	25.862	<.001
Preference Strength	5.92 (0.94)	5.33 (1.16)		
No Passage (*n* = 67)	43 (64.2%)	24 (35.8%)	5.388	.020
Preference Strength	5.28 (1.45)	4.21 (1.69)		
Negative Event				
Passage (*n* = 66)	52 (78.8%)	14 (21.2%)	21.879	<.001
Preference Strength	6.00 (1.07)	5.21 (1.42)		
No Passage (*n* = 63)	40 (63.5%)	23 (36.5%)	4.587	.032
Preference Strength	5.33 (1.44)	5.13 (1.46)		

Table 2 below summarises the descriptive data of participants' responses to the question: 'How likely do you think it is that our universe is like this?' The 'Yes' column represents the proportion of participants who reported that it is likely (5, 6, 7). The 'No' column represents the proportion of participants who reported that it is unlikely (1, 2, 3). The 'I' column represents the proportion of people who reported being indifferent between these two options. We also include the results of one-sample t-tests for each condition, which test whether the mean likelihood response is significantly different from 4. The results of these tests show that people overall think that our universe is like the one described in the passage condition. Conversely, people overall do not think that our universe is like the one described in the no-passage condition.

Table 2. Descriptive data from all conditions for participants' responses re likelihood that our universe is like the one described in the vignette.

Condition	%Yes	%No	%I	Mean	SD	t-test	*p*-value
Positive Event							
Passage	72.3	10.8	16.9	5.09	1.30	6.797	<.001
No-Passage	40.3	50.7	9.0	3.57	1.94	-1.826	.072
Negative Event							

negative past-biased preferences and 5 negative time-neutral preferences; (iv) negative event no passage condition – 11 negative past-biased preferences and 12 negative time-neutral preferences.

Passage	60.6	18.2	21.2	4.83	1.79	3.792	<.001
No-Passage	30.2	49.2	20.6	3.48	1.76	-2.364	.021

First, to confirm that there was no association between future-bias and valence, we performed a chi-square test of independence. If being a positive or negative event influences the proportion of people reporting future-bias, then we should expect a significant test result. We found *no* evidence that there is such an association χ^2 (1, N = 261) = .064, p = .800. Thus, in what follows we collapse our data across the factor valence. That is, we ignore whether the event participants were being asked to respond to was positive or negative (unless otherwise stated).

In order to test whether there was an association between future-bias and passage condition, we performed a chi-square test of independence. If the passage condition is associated with more future-bias, then there should be a significant association, with more people reporting future-bias in the passage condition than in the no-passage condition. We found evidence of such an association χ^2 (1, N = 261) = 8.612, p = .003. While most people report future-bias in both conditions, more people report future-bias in the passage condition than in the no-passage condition.

Next, to test whether the association between future-bias and passage condition differed relative to people's passage beliefs, we performed a Breslow-Day test (Breslow & Day 1980). The purpose of this test is to test the association between two variables across the levels of a third variable. In this case, does the association between future-bias and passage condition differ between participants who thought that our universe was as the vignette described and those that thought it was not or were unsure?[16] If a participant responds that it is likely that the universe in the condition is like our own, then we have some reasonably definite information about their temporal beliefs. However, if they respond that it is not like our universe (or they are unsure) then we do not have definite information about their temporal beliefs, since we do not know what it is about the metaphysics described in the condition that they disagree with. While we are interested in comparing those who believe that time robustly passes and those who believe that time is static, we cannot just compare those groups directly, since those groups differ with respect to condition, which we know influences reported future-bias. Instead, what we need to

[16] Rerunning our analyses, the reported results are not altered by removing participants who are unsure whether the universe described in the vignette is like our own. Nor are they altered by considering only those who strongly believe the vignette is like our own universe (report a 6 or 7) or strongly believe the vignette is *not* like our own universe (report a 1 or 2).

test is whether those groups differ *more* with respect to future-bias than we would expect based on condition alone. If future-bias is associated with people's temporal beliefs, then we should expect to see a significant difference in the association between future-bias and passage condition, according to whether people believe the universe in the vignette is like our own universe. However, we found no evidence of an association between people's temporal beliefs and the association between future-bias and passage condition χ^2 (1, N = 261) = .045, p = .831.[17]

Finally, to test whether there was an association between condition, temporal belief, and preference strength among participants who reported future-biased preferences, we ran a 2 x 2 between-subjects ANOVA. The two between-subject factors were condition (passage; no-passage) and temporal belief (vignette like our universe; vignette not like our universe/unsure). If people's temporal beliefs impact the strength of their future-biased preferences, then we should observe a two-way interaction between the condition and temporal belief factors. The results of this test revealed just a main effect of condition F(1, 184) = 9.387, p = .003. The strength of people's future-biased preferences was significantly greater in the passage condition (M = 5.94, SD = 1.26) than in the no-passage condition (M = 5.35, SD = 1.34).[18]

4. Discussion

Overall, then, there are two notable pairs of congruent results. First, participants in the passage condition were more likely to report future-biased preferences than were participants in the no-passage condition, and their future-biased preferences were slightly (but significantly) stronger. Second, we did not find that people who *in fact* believe that time robustly passes were more likely to be future-biased, or more strongly future-biased, than those who believe that time is static. Rather, people's reported beliefs about temporal passage had no effect on either the direction or strength of their preferences with regard

[17] Despite the considerations described in the main text, you might think it would be more natural to compare future-bias between (i) those who believe time robustly passes or do not believe time is static and (ii) those who do not believe time robustly passes or believe time is static. But in any case, if we compare these groups there is still no evidence of a significant association between temporal belief and future-bias χ^2 (1, N = 261) = .673, p = .412. While most people report future-bias, believing time robustly passes or not believing time is static is not associated with more reported future-bias than not believing time robustly passes or believing time is static.

[18] Rerunning our analyses, the reported results are not altered by removing participants unsure whether the universe described in the vignette is like our own or by including event valence as an additional factor.

to the temporal location of pleasant/unpleasant experiences. What might explain this pattern of results?

First, it could be that belief in robust passage does in fact contribute to future-bias, but that participants' actual passage beliefs had no effect in our experiment simply because our participants were extremely compliant: In endeavouring to form preferences from the perspective of our fictional astronaut, they internalized the facts about time that the astronaut was stipulated to have discovered, and their stated preferences were entirely reflective of those stipulated/simulated beliefs, rather than the participant's *actual* beliefs. On this interpretation, our results would confirm the passage belief hypothesis.

Certainly, there is good evidence that people are able to simulate things being other than they in fact take them to be. People are generally adept at imagining counterfactual scenarios, an ability that is generally accepted to underwrite our capacity to engage in causal reasoning and develop causal explanations.[19] Recent experimental work also suggests that people are able to evaluate counterfactual claims where the counterfactual world is very different indeed from the way they take the actual world to be (Roskies and Nichols 2008; Latham 2019; Latham, Miller and Norton 2020a, 2020b). This evidence, however, suggests that people's actual beliefs tend to colour their evaluations (Latham, Miller and Norton 2020a, 2020b). Similarly, research in moral psychology has found that people's personal assessments of the probabilities of outcomes in hypothetical choice scenarios can override the stipulated probabilities that they are asked to treat as features of the scenario (Ryazanov et al. 2018). And the 'curse of knowledge' phenomenon similarly illustrates the difficulty of setting aside one's actual beliefs for purposes of mental simulation.[20] So if belief in robust passage does contribute to future-bias, then while we would expect to find that simulating a belief that time does (or does not) robustly pass has an effect on future-biased preferences, we would also expect to find an interaction between this effect and people's *actual* beliefs about whether time robustly passes. We do not find any such interaction. Thus, our results seem to provide evidence *against* the passage belief hypothesis. And if we reject this hypothesis, then it seems likely that neither actual nor simulated passage beliefs explain the pattern of results seen in our experiment.

[19] For relevant work on causal reasoning, see Gopnik, Glymour, Sobel, Schulz, Kushnir, and Danks (2004), Kushnir, Gopnik, Lucas and Schultz (2010), Lagnado and Sloman (2004), Steyvers, Tenenbaum, Wagenmakers and Blum, (2003) and Sloman (2005); on causal explanation, see Einhorn and Hogarth (1986).

[20] See Nickerson (1999) for an overview of this literature.

There are, as we previously noted, versions of the temporal belief hypothesis other than the passage belief hypothesis. In particular, it might be that beliefs about temporal *ontology* (presentism vs. growing block vs. frozen block) affect future-bias. However, we do not see any plausible way in which these alternative temporal belief hypotheses could explain our results. This doesn't rule out these hypotheses, but it does suggest that there is some other factor, beyond beliefs about the metaphysics of time, that contributes to future-bias and explains our results.

A second possibility is that our results support the passage phenomenology hypothesis. Perhaps when participants imagine themselves to be the astronaut described in the vignette, the sorts of 'moving time' locutions used in the passage condition tend to produce in the participants a robust passage phenomenology. By contrast, the language used in the no-passage condition does not include any moving time locutions, and hence, so goes the thought, does not tend to produce in participants any robust passage phenomenology when they imagine being the character in the vignette. (More precisely, the vignette in the no-passage condition does contain *some* moving time locutions, but fewer than the passage condition, and all within the scope of negation.) We then find more people who are future-biased in the passage condition compared to the no-passage condition because more people in the passage condition have robust passage phenomenology, or because that phenomenology is stronger in the passage condition than in the no-passage condition. Since people can have robust passage phenomenology without believing that time robustly passes, we might expect there to be an association between people being in the passage condition (which triggers that phenomenology) and being future-biased, but no association between them believing that time robustly passes and being future-biased.[21] This is in fact what we find. This suggests that future research could profitably be directed at testing the passage phenomenology hypothesis.

A third possibility is that our results support the hypothesis that some *other* phenomenology is relevant in explaining future-biased preferences. A natural suggestion is that the use of moving time language in the passage condition produces a phenomenology that is not produced, or is produced to a lesser extent, by language used in the no-passage condition. It might be, for instance, that moving time language tends to elicit more anticipatory and retrospective experiences when compared to the language used in the no-passage condition. Suppose that were so. It is well documented that people

[21] It is natural to hypothesize, though, that people who are more disposed to experience a phenomenology as of robust passage will be more likely to develop a belief in robust passage.

experience more intense emotions during anticipation than during retrospection of the same experience. This is so for both actual and hypothetical experiences (Caruso et al. 2008; D'Argembeau and Linden 2004; Van Boven and Ashworth 2007).[22] Plausibly, if thinking about future pains evokes stronger negative emotions than thinking about past pains, this could cause us to prefer that pains be located in the past; and if thinking about future pleasures evokes stronger positive emotions than thinking about past pleasures, this could cause us to prefer that pleasures be located in the future. Call this the *emotional asymmetry hypothesis*. This hypothesis might explain the results of our experiment, if the language used in the passage condition caused participants to attend more to the imagined experience of eating their favourite or most disliked meal, and thereby heightened the emotional asymmetry between the imagined possibilities of that event being in the past or the future. Further research into the role of this kind of asymmetry would be very welcome.

One might also think that there is another category of explanations for future-bias to be considered, namely, *evolutionary* explanations. Indeed, one might even think that future-bias, insofar as it is a widespread (perhaps universal) feature of human psychology, *must* have an evolutionary explanation (or at least, that there should be a strong presumption that it does). And one might therefore think that the temporal belief hypothesis (or the temporal metaphysics hypothesis more generally) was never worth investigating in the first place. After all, the thought goes, arcane metaphysical beliefs like the belief in an objective present and robust temporal passage are not products of evolution but of culture—indeed, of a very narrow subculture, namely, academic metaphysics and philosophy of time. Any evolutionary explanation must work through much more basic, innate features of our psychology, like the drive to seek out positive experiences and avoid negative ones.[23]

While we agree that future-bias, like any widespread feature of human psychology, must have an explanation that is at least partly evolutionary (being either an adaptation, or adaptations, or a spandrel, i.e., a byproduct of adaptations), we don't think that this counts against the temporal belief hypothesis, or constitutes a distinct alternative to the hypotheses we have considered so far. On any plausible version of the temporal belief hypothesis, the beliefs in question must be nearly always implicit—basic features of how we represent the world and our relationship to it, that only become explicit in the unusual

[22] One possible explanation for this is that our episodic memories constrain our representations of the former in a way that they do not constrain our representations of the latter (Van Boven, Kane and McGraw 2009).

[23] Thanks to an anonymous referee for suggesting this line of argument.

circumstance that we are prompted to reflect consciously on the metaphysics of time. By analogy, we have many widespread implicit beliefs about other abstruse domains like physics (e.g., 'solid objects can't interpenetrate') and ethics ('it's worse to do harm than to merely allow the same harm') that are only rarely made explicit. Like these other sorts of implicit belief, the representation of events as past, present, or future, of their temporal distance from the agent, and perhaps of events or ourselves in motion through time, may be crucial to the kinds of reasoning by which humans and other cognitively sophisticated animals navigate the world.[24] Any such basic, widespread, implicit feature of the way we represent the world seems likely to have an explanation in which evolutionary forces figure prominently. This, of course, leaves open that the representation of robust passage (if indeed we have one) might be a by-product of an adaptation rather than an adaptation in its own right. Regardless, though, we don't take the temporal metaphysics hypothesis to be in competition with the claim that future-bias has an evolutionary explanation. The temporal metaphysics hypothesis is simply one way the evolutionary story might go. The emotional asymmetry hypothesis is another. And the practical irrelevance hypothesis (which we turn to shortly) is a third.[25]

Let's now turn to consider what implications, if any, our findings have for the normative status of future-bias. This is complicated, and we will limit ourselves to briefly laying out some possibilities. On the one hand, you might think that future-bias is just a matter of taste, like food preferences. This wouldn't necessarily mean that our time-biases are always exempt from rational criticism. For instance, if you always preferred apples to bananas, until forming an irrational belief that eating apples increases your risk of being struck by a meteor, then your newfound preference for bananas over apples may be irrational. But such preferences, it is natural to suppose, are *usually* exempt from rational criticism—that is, in most cases, any preference is rationally permissible. Call this *nonjudgmentalism* about future-bias. On the other hand, you might think that future-bias is fully subject to rational criticism, being rationally required if there is some feature of future experiences that makes them more significant from a prudential or other evaluative point of view than past experiences, and rationally impermissible otherwise. Call this *judgementalism* about future-bias.

[24] For discussion of the role of temporal reasoning in humans and animals, see Hoerl and McCormack (2019).

[25] Indeed, it may be that there are connections between these candidate explanations. For instance, the emotional asymmetry that we just noted might be the product of the causal irrelevance of the past: so it might be that a single underlying mechanism explains both, which in turn explains (partly) future-bias.

If we are to be judgmentalists, what kinds of considerations come into play when evaluating the rationality of our preferences? When an agent has a certain preference as a result of having certain beliefs or experiences or emotions, the *rationality* of her preferences may depend on whether she is *justified* in holding those beliefs, or taking those experiences for veridical, or regarding those emotions as fitting. For instance, if you prefer to sleep with a string of garlic around your neck because you believe that a vampire will try to bite you during the night, or thought you saw a vampire outside your window, or live in constant fear of vampires, then the rationality of your preference depends (at least partially) on whether you are justified in holding that belief, taking that experience for veridical, or regarding that fear as fitting. In such cases, where the *rationality* of a preference depends on the agent's *justifications*, we might say that the preference is *appropriate* just in case the underlying beliefs *are in fact true* or the underlying experiences *are in fact veridical* or the underlying emotions *are in fact fitting*.

Now suppose, contrary to the conclusion we have drawn from our findings, that future-bias is partly explained by belief in robust temporal passage. As has been noted, it is not clear that even if there *were* robust temporal passage, this would furnish any reason for caring more about the future than the past.[26] Still, various philosophers have taken there to be an important connection between the passage of time and the normative status of future-bias. For instance, Cockburn (1997; 1998) argues that the belief that there is no ontological difference between past, present, and future would rationally require agents to care about past, present, and future events equally. On Cockburn's view, our typical pattern of emotions toward past, present, and future events can only be rationalized if we believe in robust temporal passage. So, Cockburn thinks, unless there is some independent reason to think that future-bias is not, in fact, rational, this would be a reason to conclude that does in fact robustly pass. Similar ideas have been voiced by Prior (1959), and Pearson (2018), who straightforwardly argue that because future-bias is in fact rational, it must be the case that time robustly passes.[27]

Some of these authors think that what explains and justifies future-bias is just the

[26] For instance, Yehezkel writes: '[T]he failure to offer any substantial justification for the asymmetry in our attitudes based on the flow of time stems from the inability to offer any non-trivial account of the flow of time. It is difficult to see what difference is made by the claim that "future events are moving closer to reality," given that all that is meant by this claim is that "in the future, future events will be closer to the present." This is a mere truism, as evident by the analogous claim, regarding the past, according to which "in the past, past events were closer to the present." The attempt to justify the asymmetry between past and future based on the flow of time *per se* thus seems to collapse into triviality' (2013, pp. 6-7).

[27] Though see Maclaurin and Dyke (2002) and Suhler and Callender (2012) for critical discussion of these arguments. Thanks to an anonymous referee for drawing our attention to Cockburn's work in this area.

fact that time robustly passes, and that this justification need not go via our beliefs (see for instance Pearson (2018)). If, however, our belief that time robustly passes partly *explains* why we are future-biased (a claim that Pearson, for instance, does not make), and if time does robustly pass, and if we are justified in believing that time robustly passes, then the fact that future-bias is at least partly *explained by* belief in robust passage would seem to count in favour of the conclusion that future-bias is appropriate (rational). And inversely, if either of the conditions above are false (regarding the truth/justification of belief in robust temporal passage, or the justificatory link from passage to future-bias), then the fact that future-bias is at least partly explained by belief in robust passage would count against the conclusion that future-bias is appropriate/rational.

The same analysis applies, *mutatis mutandis*, to the passage phenomenology and emotional asymmetry hypotheses.

Under any of these hypotheses, the evaluative stakes are lower for nonjudgementalists, who may simply hold that future-bias is permissible (but not required) regardless of how we evaluate the beliefs/phenomenology/emotions that underlie it. But a nonjudgementalist might still think that even if both future-bias and time-neutrality are *potentially* appropriate, our *actual* future-biased preferences are inappropriate if they are the product of false beliefs, or non-veridical phenomenology, or unfitting emotions.

On this last point, though, it is especially unclear whether nonjudgmentalists should care whether future-bias is the result of a non-veridical phenomenology or an unfitting pattern of emotions. Suppose that, on your way to work, you have a non-veridical experience of an apple (e.g., you hallucinate an apple), and that because the illusory apple looks so delicious, you develop a preference for apples which explains why you choose an apple over a banana at lunchtime. Or suppose that you were told a scary story about bananas as a child and came to develop an unfitting fear of bananas as an adult, which explains your preference for apples over bananas at lunchtime. Most of us, being nonjudgmentalists about food preferences, may be inclined to say that the suspect aetiology of your lunchtime preference does not impugn its appropriateness or rationality. And one might very well say the same for our future-biased preferences, if they turn out to be explained by nonveridical passage phenomenology and/or an unfitting emotional asymmetry between anticipation and retrospection.

Additionally, it is worth noting that our findings are compatible with the conclusion that future-bias is explained by *several* factors. This suggests that evaluating the appropriateness/rationality of future-bias might be significantly more complicated than

philosophers have hitherto supposed. For it could well be that any particular future-biased preference is partly determined by a number of factors, some of which count in favour of and some of which count against its appropriateness/rationality. Indeed, there is some reason to suppose that this is so. Previous empirical work (Latham et al. 2020) suggests that future-bias is mitigated when we consider scenarios where the past is causally accessible: that is, it supports the practical irrelevance hypothesis.

Suppose that future-biased preferences are partly explained by the causal irrelevance of the past, and partly by other factors, like non-veridical passage phenomenology and/or an unfitting anticipation/retrospection asymmetry. Then it could be that these different factors tend to pull in difference directions when it comes to the normative status of the resulting preferences.

For instance, one might think that the practical irrelevance of the past is a reason to attach less evaluative weight to past events. At the very least, one might think that future-biased preferences are, in such circumstances, permissible, since they are harmless: where the past is causally irrelevant, future-biased preferences cannot guide action, and so, one might think, they cannot lead us into actions that leave us worse off. If one thought this, then determining the overall normative status of these preferences would not be entirely straightforward.

Of course one might not think that the practical irrelevance of the past is a reason to attach less evaluative weight to past events. The idea that future-biased preferences might be harmless has come under considerable scrutiny. This assumption is less straightforward than it might appear: an agent who is risk-averse (Dougherty 2011), regret-averse (Greene and Sullivan 2015), or an evidential decision theorist (Tarsney 2017) may choose differently because she is future-biased, even when she cannot affect the past, potentially in ways that leave her worse off overall. (For instance, Dougherty (2011) argues that agents who are risk-averse and future-biased can be made into 'pain pumps'.)

Our general point, here, is just that if these preferences are explained by multiple factors, then it may be that these factors tend to confer a different normative status on the resulting preferences, and hence determining their overall status might be complicated.

Finally, we should acknowledge one important limitation of our study: It is hard to be certain that our vignettes successfully conveyed the ideas of robust passage and static time to our participants. The disagreement between the A- and B-theories (in particular, the difference between robust passage and mere B-theoretic "succession" of earlier events by later events) is hard to state clearly and without reliance on metaphorical/figurative

language to the point that many philosophers find the difference obscure, and even doubt whether it is genuine (e.g., Callender 2000).[28] However, the fact that most participants were able to form a judgment that time as described in their vignette either did or didn't resemble time in the real world, and that (in-line with previous research (Latham, Miller and Norton 2019, 2020a) most participants believed the passage vignettes resembled the actual world and did not believe that the no-passage vignettes resembled the actual world (see Table 2) is some evidence that we were effective in conveying the idea of robust temporal passage. Nevertheless, there would be value in other tests of the passage belief hypothesis that convey the idea of/probe belief in robust passage in other ways.

5. Conclusion

Much more empirical work is required before we have a complete picture of what explains our future-biased preferences. The current study sheds light on some of the factors that have been thought to influence these preferences. Our results suggest that, contrary to what many philosophers have supposed, beliefs about robust temporal passage probably are not a major source of future-bias. Nevertheless, our results do clearly suggest that there is a factor (or set of factors) that is in some way connected to robust temporal passage and plays a role in future-bias. We have suggested that this factor might either be our phenomenology as of robust temporal passage, or the asymmetric emotions associated with anticipation and retrospection. Further research can help to determine which, if either, of these hypotheses is correct and, in turn, inform philosophical debates about the implications and the normative status of future-bias.

References

Ahler, D., Roush, C., & Sood, G. The micro-task market for lemons: data quality on Amazon's Mechanical Turk. Unpublished manuscript, 22 January 2020. URL: http://gsood.com/research/papers/turk.pdf

Boroditsky, L. (2000). 'Metaphoric Structuring: Understanding Time though Spatial Metaphors'. *Cognition* 75(1): 1–28.

Brink, D. O. (2011). 'Prospects for Temporal Neutrality'. in Craig Callender (ed.) T*he Oxford Handbook of Philosophy of Time*. Oxford: Oxford University Press.

Callender, C. (2000). 'Shedding Light on Time' *Philosophy of Science* 67: S587–S599.

[28] Of course, even if the distinction between robust passage and static time (and therefore belief in one as opposed to the other) is ultimately incoherent, that does not rule out the passage belief hypothesis – incoherent beliefs can still be psychologically consequential.

Caruso, E., D. Gilbert and T. Wilson. (2008). 'A wrinkle in time: asymmetric valuation of past and future events.' *Psychological Science* 19: 796–801.

Clark, H. (1973). 'Space, Time Semantics and the Child' *Cognitive development and the acquisition of language*. Edited by T. E. Moore NY Press. pp 28–63.

Cockburn, D. (1997). *Other Times: Philosophical Perspectives on Past, Present, and Future*. Cambridge: Cambridge University Press.

Cockburn, D. (1998). 'Tense and Emotion.' In R. Le Poidevin (ed.) *Questions of Time and Tense*. Oxford: Clarendon Press: 77–91.

Craig, W.L. (1999). 'Tensed time and our differential experience of the past and future.' *Southern Journal of Philosophy* 37: 515–37.

D'Argembeau, A. and Linden, M. V (2004) 'Phenomenal characteristics associated with projecting oneself back into the past and forward into the future: Influence of valence and temporal distance'. *Consciousness and Cognition* 13: 844–858

Dorsey, D. (2018). 'Prudence and Past Selves.' *Philosophical Studies* 175(8):1901–1925.

Dougherty, T. (2011). On whether to prefer pain to pass. *Ethics* 121: 521–37.

Dougherty, T. (2015). 'Future-Bias and Practical Reason'. *Philosophers' Imprint* 15.

Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, *99*(1).

Frijda, N. H. (1988). 'The laws of emotion'. *American Psychologist* 43: 349–358

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). 'A theory of causal learning in children: Causal maps and Bayes nets.' *Psychological Review,* 111:1–30.

Greene, P. and M. Sullivan. (2015). 'Against time bias.' *Ethics* 125: 947–70.

Greene, P., Latham, A. J., Miller, K and Norton, J. (2020). 'Hedonic and non-hedonic bias towards the future'. *Australasian Journal of Philosophy* https://doi.org/10.1080/00048402.2019.1703017

Hare, C. (2007). 'Self-bias, time-bias, and the metaphysics of self and time.' *Journal of Philosophy* 104: 350–73.

Hare, C. (2013). 'Time – the emotional asymmetry.' In *A Companion to the Philosophy of Time*, ed. A. Bardon and H. Dyke, 507–20. Hoboken, NJ: Wiley-Blackwell.

Heathwood, C. (2008). 'Fitting Attitudes and Welfare' *Oxford Studies in Metaethics* 3: 47–73.

Hoerl, C. & McCormack, T. (2019). "Thinking in and about time: A dual systems perspective on temporal cognition.' *Behavioural and Brain Sciences* 52(244): 1-17.

Horwich, P. (1987). *Asymmetries in Time: Problems in the Philosophy of Science*. Cambridge, MA: MIT Press.

Hume, D. 2000 (1739). *A Treatise of Human Nature*. Oxford: Oxford University Press.

Kauppinen, A. (2018). 'Agency, Experience, and Future Bias.' *Thought: A Journal of Philosophy* 7(4): 237–245.

Kushnir, T., Gopnik, A., Lucas, C. & Schulz, L. (2010). 'Inferring Hidden Causal Structure.' *Cognitive Science,* 34(1): 148–160

Lagnado, D. A., & Sloman, S. A. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 856–876.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.

Latham, A. J. (2019). 'Indirect Compatibilism' PhD Dissertation. https://philpapers.org/rec/LATIC

Latham, A. J., Miller, K., and Norton, J. (2019). 'Is our naïve theory of time dynamical?' *Synthese*. DOI: 10.1007/s11229-019-02340-4[2]

Latham, A. J., Miller, K and Norton, J. (2020a) 'Do the Folk Represent Time as Essentially Dynamical?' *Inquiry*. https://doi.org/10.1080/0020174X.2020.1827027

Latham, A.J., Miller, K. and Norton, J. (2020b). 'An empirical investigation of the role of direction in our concept of time'. *Acta Analytica.* DOI: 10.1007/s12136-020-00435-z

Latham, A.J., Miller, K. and Norton, J. (2020c). 'An empirical investigation of the role of direction in our concept of time'. *Acta Analytica.* DOI: 10.1007/s12136-020-00435-z

MacBeath, M. (1983). 'Mellor's Emeritus headache.' *Ratio*, 25:81–88.

Maclaurin, J. & Dyke, H. (2002). "Thank goodness that's over': the evolutionary story.' *Ratio* 15:276–92.

McTaggart, J. E. (1908). 'The unreality of time.' *Mind* 17(68):457–474.

Mellor, D. H. (1981). 'Thank goodness that's over.' *Ratio*, 23:20–30.

Mellor, D. H. (1983). 'MacBeath's soluble aspirin.' *Ratio*, 25:89–92.

Nickerson, R. S. (1999). How We Know—And Sometimes Misjudge—What Others Know: Imputing One's Own Knowledge to Others. *Psychological Bulletin* 125: 737–759.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Pearson, O. (2018). 'Appropriate Emotions and the Metaphysics of Time.' *Philosophical Studies* 175(8):1945–1961.

Prior, A.N. (1959). 'Thank goodness that's over.' *Philosophy* 34:12–7.

Roskies, A and S Nichols, (2008). 'Bringing Moral Responsibility Down to Earth.' *Journal of Philosophy* 105 (7): 371-388.

Ryazanov, A. A., Knutzen, J., Rickless, S. C., Christenfeld, N. J., & Nelkin, D. K. (2018). 'Intuitive probabilities and the limitation of moral imagination'. *Cognitive Science*, *42*, 38–68.

Schlesinger, G. (1976). 'The stillness of time and philosophical equanimity.' *Philosophical Studies* 30:145–59.

Sloman, S. A. (2005). *Causal Models: how people think about the world and its alternatives*. Oxford: OUP.

Smith, Q. (1994). 'The Phenomenology of A-Time,' in L. Nathan Oaklander and Quentin Smith, eds., *The New Theory of Time*. New Haven: Yale University Press. pp. 351–59.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive science*, *27*(3), 453–489.

Suhler, C. and C. Callender. (2012). 'Thank goodness that argument is over: explaining the temporal value asymmetry.' *Philosophers' Imprint* 12:1–16.

Sullivan, M. (2018). *Time Biases*. Oxford: Oxford University Press.

Tarsney, C. (2017). "Thank goodness that's Newcomb': the practical relevance of the temporal value asymmetry". *Analysis* 77(4): 750–9.

van Boven, L. and Ashworth , L. (2007). 'Looking Forward, Looking Back: Anticipation Is More Evocative Than Retrospection' *Journal of Experimental Psychology: General* 136(2): 289–300.

van Boven, L., Kane, J., and McGraw, A. P. (2009). 'Temporally asymmetric constraints on mental simulation: Retrospection is more constrained than prospection'. In K. D. Markman, W. M. P. Klein, & J. A. Suhr (Eds.), *Handbook of imagination and mental simulation* (pp. 131–147). New York, NY: Psychology Press.

Yehezkel, G. (2014). 'Theories of Time and the Asymmetry in Human Attitudes.' *Ratio* 27(1):68–83.

Zimmerman, D. (2005). 'The A-Theory of Time, The B-Theory of Time, and 'Taking Tense Seriously." *Dialectica* 59(4): 401–457.