Journal of the American Philosophical Association (2024) 588–614
DOI:10.1017/apa.2023.14

Exploring Arbitrariness Objections to Time Biases

ABSTRACT: *There are two kinds of time bias: near bias and future bias. While philosophers typically hold that near bias is rationally impermissible, many hold that future bias is rationally permissible. Call this* normative hybridism. *According to arbitrariness objections, certain patterns of preference are rationally impermissible because they are arbitrary. While arbitrariness objections have been leveled against both near bias and future bias, the kind of arbitrariness in question has been different. In this article we investigate whether there are forms of arbitrariness that are common to both kinds of preferences and, hence, whether there are versions of the arbitrariness objection that are objections to both near bias and future bias. If there are, then this might go some way toward undermining normative hybridism and to defending thoroughgoing time-neutralism.*

KEYWORDS: time bias, future bias, time neutrality, prudence, temporal experience, temporal emotions

Introduction

We mortals have time biases. (Although the term 'bias' may have negative connotations, here we use it in a purely descriptive way.) A person is biased toward the future (i.e., 'future biased') when, all else being equal, they prefer positively valenced events to be future rather than past or prefer negatively valenced events to be past rather than future. Imagine waking up in a hospital bed, dazed and amnesic, knowing only that one of the following two conditions obtains: you are about to have a painful operation, or you have just had an equally painful operation. Research shows you probably prefer the operation to be past (Caruso, Gilbert, and Wilson 2008; Greene et al. 2020). In fact, you probably still prefer the past operation even if it is much more painful than the future one (Lee et al. 2020; Greene et al. 2021a).

On the other hand, a person is biased toward the near (i.e., 'near biased') when, all else being equal, they prefer positive events to be near rather than distant, or prefer negative events to be distant rather than near. Suppose you are offered some dessert you like, which you can either eat now or tomorrow. If, all else being equal, you prefer to eat it now, you are near biased.[1] Near bias can be prospective or

[1] For instance, Thaler (1981) showed that people prefer less money given now to more money given later, and Hausman (1979) found that people were willing to buy cheaper air conditioners with higher operating costs down

retrospective. We shall primarily focus on the former—that is, on the preference for positive events to be in the near rather than far future and for negative events to be in the far rather than near future – as prospective near bias has been more empirically investigated.

In each case the 'all else being equal' qualification is important. Suppose you prefer to eat the dessert now because (say) you happen to be very hungry now, and you do not expect yourself to be equally hungry tomorrow, or you are worried that the dessert might go bad by tomorrow: then you are merely apparently near biased. Likewise, if you prefer the past operation because (say) you believe the sooner the operation is done, the likelier your illness will be cured, then you are merely apparently future biased. On the contrary, you are genuinely future biased or near biased only when your preferences are sensitive to the temporal locations of the events themselves as opposed to variations in uncertainty, intrinsic values, and so on, which result from variations in temporal locations (see also Lowry and Peterson 2011: 490). In what follows, we will focus on genuine time biases and hence drop the 'genuine' locution except where qualification is necessary. (However, in section 4, we will question whether our experiments, though intended to uncover our patterns of genuine time biases, in fact track merely apparent time biases.)

Time biases come in various strengths. It is helpful to think of both near bias and future bias as forms of discounting subjective (expected) utility along the temporal dimension, so that situations of 'equal payoff' and 'unequal payoff' (prior to discounting) can be united into one model. The aforementioned case of past or future operations, both equally painful, is one of equal payoff, whereas an alternative version of the case (Parfit 1984: 165–66), in which the past operation is more painful than the future one, is one of unequal payoff. If I am only slightly future biased toward surgical pain, I might prefer the past operation in the original case but not in the alternative version. However, if I am strongly future biased, or even absolutely future biased (if I do not regard past pain as of any value at all), then I might well prefer the past operation in both scenarios.

Studies in psychology, behavioral economics, and philosophy have found considerable intrapersonal and interpersonal variation in the presence or absence of future and near bias, as well as their strengths (see Frederick, Loewenstein, and O'Donoghue 2002). Notably, for the present purposes, both future and near bias have been shown to be sensitive to the type of goods/events in question (Frederick, Loewenstein, and O'Donoghue 2002; Greene et al. 2020), and future bias has been shown to be sensitive to the valence (positive or negative) of the goods/events in question (Greene et al. 2020, 2021a, 2022a).

At least some of this variation has been fodder for philosophical arguments regarding the normative status of these time biases. We will call arguments that

the line. For overviews see Soman et al. (2014), Frederick, Loewenstein, and O'Donoghue (2002), Ainslie and Haslam (1992) and Hardisty and Weber (2020). By contrast, retrospective near bias has largely been overlooked, but see Yi, Gatchalian, and Bickel (2006), Bickel et al. (2008), and Greene et al. (2021b) on the correlation of prospective and retrospective near biases.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

appeal to factors such as these *arbitrariness arguments*. These arguments seek to show that the preferences in question are arbitrary and hence rationally impermissible. These kinds of arguments can be distinguished from what we might call *upshot arguments*, which seek to show that the preferences are ones that can make agents worse off in some way.[2]

There is, of course, a perfectly general arbitrariness argument against both near bias and future bias: namely, that it is arbitrary to discount the expected utility of some person-stage simply in virtue of where in time that stage is located. Temporal location, one might think, is not normatively significant, and hence any such discounting is arbitrary. Arguments of this kind against near bias can be traced back at least as far as Sidgwick (1884) and have been taken up by many authors (Rawls 1971: 293–94; Parfit 1984: 124; Lowry and Peterson 2011: 493). Imagine someone who is otherwise normal but is indifferent to pleasures or pains on future Tuesdays while knowing perfectly well that there is nothing special about Tuesdays (Parfit 1984: 124). It seems that this preference is groundless, arbitrary, and hence irrational. Near bias, in virtue of giving more weight to temporal locations that are merely closer, seems arbitrary in much the same way.

While the general arbitrariness argument has been persuasive when it comes to near bias, it has been much less persuasive in the case of future bias. That is because many philosophers have thought that there is some important asymmetry —for example, a metaphysical one (Prior 1959; Schlesinger 1976; Craig 1999; Pearson 2018)—between past and future, which means that future-biased preferences are not arbitrary at all. Hence the general arbitrariness argument fails when it comes to future bias. Accordingly, while near bias is not rationally permissible, future bias is rationally permissible (Heathwood 2008; Hedden 2015) or even obligatory (Prior 1959; Schlesinger 1976; Craig 1999; Pearson 2018). We call this view *normative hybridism*.

This orthodoxy has recently been challenged by time-neutralists who present a series of what Greene et al. (2022c) call preference-pattern arguments. In essence, these are arguments that try to show that future-biased preferences are arbitrary not because they are sensitive to temporal location, but because they are sensitive to some other normatively irrelevant features. Several such preference-pattern arguments have been offered in the literature, including the argument from first-person versus third-person preferences (Parfit 1984: 181, Dougherty 2015: 3) and the argument from hedonic versus non-hedonic preferences (Brink 2011; Dougherty 2015: 3). Each of these arguments tries to show that future-biased preferences are sensitive to some factor that is normatively irrelevant, indicating that they are caused by some non-reason-tracking mechanisms, such as cognitive or affective biases, and are therefore arbitrary and irrational.

[2] For arguments of this kind, see Trout (2007) regarding near bias. When it comes to future bias, it has been suggested that although we cannot change the past, future bias in combination with other rational principles or tendencies can guide *future* actions and make people worse off overall (Dougherty 2015). In response, Kauppinen (2018) argues that future bias is irrational only when it is action-guiding.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Interestingly, though, although some research suggests that there is a connection between near bias and future bias (Latham, Miller, and Norton 2023), there has been no attempt to see whether there might be some normatively irrelevant feature to which both kinds of preferences are sensitive. If there were, this would allow time-neutralists to present a new preference-pattern argument. This argument would, like the general arbitrariness argument, target both near-biased and future-biased preferences but without appealing to the idea that being sensitive to temporal location is always arbitrary.

In this article we aim to determine whether either of two factors that are, very plausibly, normatively irrelevant—*valence* and *kind of event*—are ones to which near-biased and future-biased preferences are both sensitive. If they are, this would permit time-neutralists to offer one (or two) new preference-pattern arguments that cite valence or event kind as being a common factor to which near-biased and future-biased preferences are sensitive, such that being sensitive to this factor renders the resulting preferences arbitrary.

In section 1 we discuss current relevant research, before outlining our methodology and results in section 2. In section 3 we discuss the implications of these results for arbitrariness objections to near and future bias and, in turn, the implications for arguments against normative hybridism.

1. Near Bias and Future Bias

While the general arbitrariness argument has often been resisted by philosophers, it has typically been assumed that a version of the argument succeeds when it comes to near bias (Rawls 1971: 293–94; Parfit 1984: 124; Lowry and Peterson 2011: 493). Philosophers tend to assume that it is normatively irrelevant where in the future (near or far) states of affairs are located, and hence conclude that near bias (prospective near bias at least) is arbitrary and objectionable. Although it has long been observed that near bias does not have a uniform pattern—rather, the discount rates for different types of goods/events vary considerably, both interpersonally and intrapersonally (Frederick, Loewenstein, and O'Donoghue 2002)—philosophers have rarely argued against the rational permissibility of near bias by appealing to preference-pattern arguments that cite the wide variation in discount rates across goods.

This is likely because they take it that such arguments are redundant. (Economists argue for the rational impermissibility of certain forms of discounting, namely, discounting that is hyperbolic. This is not on the grounds of arbitrariness, but because such discounting leads to people being susceptible to money pumping.) By contrast, appeals to preference-pattern arguments have been pursued when it comes to evaluating the normative status of future bias. Time-neutralists have used preference-pattern arguments to try to show that there are various normatively irrelevant features to which future-biased preferences are sensitive and, hence, that such preferences are objectionably arbitrary. Here is a general schema of a preference-pattern argument:

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Preference-Pattern Argument

1. Preferences that are sensitive to normatively irrelevant factors are objectionably arbitrary.
2. We have reason to avoid having objectionably arbitrary preferences.
3. Future-biased preferences are sensitive to *F*.
4. *F* is normatively irrelevant.
5. Therefore, future-biased preferences are sensitive to a normatively irrelevant factor (from 3, 4).
6. Therefore, we have reason to avoid having future-biased preferences (from 1, 2, 5).

Time-neutralists have offered various candidates for factor *F*. For instance, it has been thought that people are only future biased when it comes to hedonic events and not non-hedonic[3] ones (Brink 2011; Dougherty 2015: 3); thus, various time-neutralists have substituted for *F* 'whether a preference is hedonic or non-hedonic'. It has also been thought that people have future-biased preferences only in first-person conditions and not in third-person conditions (Parfit 1984: 181; Dougherty 2015: 3; cf. Hare 2013). Accordingly, time-neutralists have substituted for *F* 'whether the preference is first- or third-person' (Brink 2011: 378–79; Dougherty 2015: 3).

Unfortunately for time-neutralists, however, these predictions have not been borne out by empirical investigation. Recent empirical work has failed to find any asymmetry between first-person and third-person preferences (Greene et al. 2020, 2021a) or between hedonic and non-hedonic preferences (Greene et al. 2020). These particular examples of preference-pattern arguments appear to be unsound.

Nevertheless, time-neutralists can offer preference-pattern arguments that appeal to other factors. Greene et al. (2021a, 2022a) recently found that people are sensitive to the valence of events: they are more future biased when it comes to negative events than positive ones. Call this the valence factor. We also know that when it comes to near bias, people are sensitive to the kind of good/event in question. Call this the event factor.

Recent research has found a moderate association between near-biased and future-biased preferences (Latham, Miller, and Norton 2023), suggesting that perhaps there is a factor to which *both* are sensitive. This would offer the possibility of filling out a more general preference-pattern argument such as the one below:

[3] There is no general agreement on how exactly the hedonic/non-hedonic distinction should be drawn (see, for example, Labukt 2012; Bramble 2016), and we shall not attempt to offer one here. Suffice it to say that paradigmatic hedonic events include pure pleasure and pain, valenced emotions, and pleasant and unpleasant sensations, while paradigmatic non-hedonic events include winning a lottery, losing a job, and being betrayed by a friend.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Generalized Preference-Pattern Argument

7. Preferences that are sensitive to normatively irrelevant factors are objectionably arbitrary.
8. We have reason to avoid having objectionably arbitrary preferences.
9. Future-biased and near-biased preferences are both sensitive to *F*.
10. *F* is normatively irrelevant.
11. Therefore, future-biased and near-biased preferences are both sensitive to a normatively irrelevant factor (from 9, 10).
12. Therefore, we have reason to avoid having both future-biased and near-biased preferences (from 7, 8, 11).

The success of such an argument would go some way toward undermining normative hybridism. This article empirically investigates two candidate factors to substitute for *F*: valence and kind of event

We developed three broad hypotheses. The first is that people's near-biased and future-biased preferences will be sensitive to the event factor. Events can vary in all manners of ways, and here we focus on two dimensions of variation that have not yet been examined. First, we focus on the distinction between sensations and moods. Second, we focus on the distinction between kinds of sensations. Our first experiment tests the kind of hedonic event hypothesis. According to this hypothesis, both near-biased and future-biased preferences are sensitive to whether the hedonic event in question is a sensation or a mood. Our second experiment tests the kind of sensation hypothesis. According to this hypothesis, both near-biased and future-biased preferences are sensitive to the kind of sensation (touch, taste, smell, sight).

Our second broad hypothesis is that people's near-biased and future-biased preferences will be sensitive to the valence factor. Inspired by previous research showing that people are more future biased about negative events than positive events (Greene et al. 2021a, 2022a), we tested the negative valence hypothesis, according to which people will be both more future biased and more near biased with regard to negative events than positive ones. (Talk of sensitivity to valence comes from Greene et al. [2021a] and Greene et al. [2022c]. In psychology this is also known as the sign effect, which refers to the phenomenon that 'people discount future positives more than future negatives' [Molouki, Hardisty, and Caruso 2019: 1674].) Both our first and second experiments test the negative valence hypothesis.

Our third broad hypothesis, made based on prior research by Latham et al. (2023), is that we would find an association between future-biased and near-biased preferences. Thus, we made the following specific predictions (experimental data, hypotheses, and materials can be found at https://osf.io/pjznb/).

H1: [Experiment 1] There will be differences in future bias across different hedonic events of sensation and mood (kind of hedonic event hypothesis).
H2: [Experiment 1] There will be differences in near bias across different hedonic events of sensation and mood (kind of hedonic event hypothesis).

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

H3: [Experiment 2] There will be differences in future bias across different kinds of sensation (kind of sensation hypothesis).
H4: [Experiment 2] There will be differences in near bias across different kinds of sensation (kind of sensation hypothesis).
H5: [Experiments 1 and 2] There will be an association between future-biased and near-biased preferences (association hypothesis).
H6: [Experiments 1 and 2] People will be more future biased about negative events than positive ones (negative valence hypothesis).
H7: [Experiments 1 and 2] People will be more near biased about negative events than positive ones (negative valence hypothesis).

2. Methodology and Results

2.1 Experiment 1 Methodology

2.1.1 Participants. In total, 583 people participated in the study. Participants were US residents, recruited and tested online using Amazon Mechanical Turk, and compensated $1 for their time. MTurk participants had a Human Intelligence Task (HIT) approval rate of at least 95 percent and had their HITs approved at least 1,000 times. In total, 342 participants were excluded for failing either to follow instructions or to answer *all* the attentional check and comprehension questions correctly. The remaining sample was composed of 241 participants (103 female, 4 trans/nonbinary; mean age 41.97; SD = 12.25; range 21–75). Ethics approval for these studies was obtained from the University of Sydney Human Research Ethics Committee. Informed consent was obtained from all participants prior to testing. The survey was conducted online using Qualtrics.

2.1.2 Materials and Procedure. Participants were randomly assigned to one of four conditions, which included every combination of valence (positive vs. negative) and kind of hedonic event (sensation vs. mood). Because the positive and negative vignettes differ only minimally, we can present them together:

> **Experiment 1:**
> **Positive/Negative Sensation/Mood**
> Imagine that three months ago you had a genetic test, and the results showed that you are very likely to develop a fatal disease in ten years. Luckily, three months ago the doctor gave you a pill that prevents this disease from developing. You took the pill in his office, and so you will not develop that disease.
> The pill is very safe and is certain to have no long-term side effects. The medication does, however, have one short-term side effect. At some time during the 12 months after you have ingested the pill, it causes the brain to misinterpret certain signals, and as a result causes three consecutive days of intense {[pain]/[pleasure]}/{[euphoria and happiness]/[depression and sadness]} after which these side effects cease, and you return to normal.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

You wake up one morning after a restless night and for a moment cannot remember whether you have already experienced these side effects.

After reading the vignette, participants responded to four comprehension questions and were offered a choice of (a) True or (b) False.

In this vignette you were asked to imagine that:

(a) Three months ago you had a genetic test, which shows you are likely to develop a fatal disease in 10 years time.
(b) Having taken the pill, you will avoid developing the fatal disease.
(c) You wake up one morning and remember that you already experienced the pill's side effects yesterday.
(d) The pill will cause you to experience three consecutive days of high fever.

Participants who failed to answer these questions correctly were excluded from the analysis. Participants then saw two sets of questions, one probing whether and to what extent they have prospective near-biased preferences and one probing whether and to what extent they have future-biased preferences. The order in which they saw these questions was randomized.

Participants were asked to 'Please indicate your preference using one of the following statements':

(a) I would prefer to learn that I will experience the side effects of the pill tomorrow, and not in five months' time.
(b) I would prefer to learn that I will experience the side effects of the pill in five months' time, and not tomorrow.
(c) I have no preference between these options.

Participants were asked to 'Please indicate your preference using one of the following statements':

(a) I would prefer to learn that I will experience the side effects of the pill tomorrow and did not experience the side effects three days ago.
(b) I would prefer to learn that I experienced the side effects of the pill three days ago and will not experience the side effects tomorrow.
(c) I have no preference between these options.

2.2 Experiment 2 Methodology

2.2.1 *Participants.* In total, 1,171 people participated in the study. Participants were US residents, recruited and tested online using Amazon Mechanical Turk, and compensated $1 for their time. Again, we used those MTurk participants who have a HIT approval rate of at least 95 percent and who have had their HITs approved at least 1,000 times. As many as 668 participants had to be excluded for failing to answer the questions or failing one of the attentional check or

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

comprehension questions. The remaining sample was composed of 503 participants (212 female, 4 trans/nonbinary; mean age 41.16; SD = 11.79; range 19–76). Ethics approval for these studies was obtained from the University of Sydney Human Research Ethics Committee. Informed consent was obtained from all participants prior to testing. The survey was conducted online using Qualtrics.

2.2.2 *Materials and Procedure.* Participants were randomly assigned to one of eight conditions, which include every combination of valence (positive vs. negative) and kind of hedonic event (smell vs. taste vs. touch vs. sight). Because the positive and negative vignettes differ only minimally, we can present them together:

> **Positive/Negative Smell/Taste/Touch/Sight**
> Imagine that three months ago you had a genetic test and the results showed that you are very likely to develop a fatal disease in ten years. Luckily, three months ago the doctor gave you a pill that prevents this disease from developing. You took the pill in his office, and so you will not develop that disease.
> The pill is very safe and is certain to have no long-term side effects. The medication does, however, have one short-term side effect. At some time during the 12 months after you have ingested the pill, it causes the brain to misinterpret certain signals and as a result causes three consecutive days in which you experience a {persistent, strong, extremely (un) pleasant taste in your mouth}/{persistent, strong, extremely (un) pleasant smell}/{persistent, strong, extremely (un)pleasant sensation on your skin}/{persistent, strong, extremely (un)pleasant visual effect} after which these side effects cease, and you return to normal.
> You wake up one morning after a restless night and for a moment cannot remember whether you have already experienced these side effects of the pill.

After reading the vignette, participants responded to the same comprehension questions as in experiment 1. Participants who failed to answer these questions correctly were excluded from the analyses. Participants then saw the same set of two probe questions as they saw in experiment 1, one probing whether and to what extent they have prospective near-biased preferences and one probing whether and to what extent they have future-biased preferences.

2.2.3 *Results*

2.2.3.1 *Experiment 1 Results.* Before reporting the statistics, we will first summarize our major findings with respect to our hypotheses. First, consider the kind of hedonic event hypothesis according to which there would be differences in future bias (*H1*) and near bias (*H2*) across different hedonic events of sensation and mood. Neither hypothesis was supported. We found that people's future-biased preferences and near-biased preferences were not significantly different across different hedonic events of sensation and mood. Next, consider the negative valence hypothesis according to which people will be more future

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

biased (*H6*) and more near biased (*H7*) about negative events than about positive ones. We found evidence for *H6*: more people had future-biased preferences when it came to negative events compared to positive ones. In contrast, *H7* was not supported. Contrary to what we predicted, we found that when it came to negative events more people had *far*-biased rather than near-biased preferences. (Far bias is the converse of near bias—i.e., preferring positive [/negative] events to be temporally distant [/close] rather than close [/distant].) Finally, consider the association hypothesis (*H5*) according to which there will be an association between future-biased and near-biased preferences. This hypothesis was not supported. We did not find robust evidence of an association between people's future-biased and near-biased preferences.

Table 1 presents descriptive data of participants' responses regarding their future-biased preferences across all conditions. The 'FB' column represents the number of participants who reported future-biased preferences. The 'PB' column represents the number of participants who reported *past*-biased preferences. (Past bias is the converse of future bias—i.e., preferring positive [/negative] events to be past [/future] rather than future [/past].) The 'NP' column represents the number of participants who reported time-neutral preferences.

To investigate whether there were any differences in future-biased preferences across the conditions, we ran a chi-square test of homogeneity. This test showed that there was a significant difference in people's future-biased preferences across our conditions: χ^2 (6, N = 241) = 15.281, p = .018.

To identify the source of this effect, we ran separate chi-square tests of homogeneity for valence (positive; negative) and kind (sensation; mood). The results of these tests found a significant effect of valence, χ^2 (2, N = 241) = 8.800, p = .012, and no significant effect of kind, χ^2 (2, N = 241) = 4.455, p = .108, on people's future-biased preferences. Post-hoc comparisons with a Bonferroni correction showed that people differed in their future-biased preferences in negative conditions. Standardized residuals showed that *more* people reported having a future-biased preference (z = 4.622, p < .001) and *fewer* people reported having a time-neutral preference (z = -3.971, p < .001) in negative conditions.

Table 2 summarizes the descriptive data of participants' responses regarding their near-biased preferences across all conditions. The 'NB' column represents the number of participants who reported near-biased preferences. The 'FrB' column represents the number of participants who reported far-biased preferences. The 'NP' column represents the number of participants who reported a time-neutral preference.

Table 1. Descriptive data of participants' responses to the future-biased prompt in Experiment 1

Valence	Kind	FB	PB	NP
Positive	Sensation	28 (46.7%)	15 (25.0%)	17 (28.3%)
	Mood	24 (38.1%)	16 (25.4%)	23 (36.5%)
Negative	Sensation	35 (63.6%)	12 (21.8%)	8 (14.5%)
	Mood	28 (44.4%)	24 (38.1%)	11 (17.5%)

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Table 2. Descriptive data of participants' responses to the near-biased prompt in Experiment 1

Valence	Kind	NB	FrB	NP
Positive	Sensation	25 (41.7%)	15 (25.0%)	20 (33.3%)
	Mood	30 (47.6%)	11 (17.5%)	22 (34.9%)
Negative	Sensation	17 (30.9%)	29 (52.7%)	9 (16.4%)
	Mood	13 (20.6%)	39 (61.9%)	11 (17.5%)

To investigate whether there was any difference of near-biased preferences across the conditions, we ran a chi-square test of homogeneity. The test showed that there was a significant difference in near-biased preferences across our conditions: χ^2 (6, N = 241) = 36.150, p < .001.

Once again, to identify the source of this effect, we ran separate chi-square tests of homogeneity for valence and kind. The results of these tests found a significant effect of valence, χ^2 (2, N = 241) = 33.836, p < .001, and no significant effect of kind, χ^2 (2, N = 241) = 0.151, p = .927. Post-hoc comparisons with a Bonferroni correction showed that people differed in their near-biased preferences in both negative and positive conditions. Standardized residuals showed that *more* people reported having a far-biased preference (z = 5.598, p < .001), and *fewer* people reported having a time-neutral preference (z = -3.775, p < .001) in negative conditions. In contrast, *fewer* people reported having a far-biased preference (z = -2.869, p = .004), and *more* people reported having a near-biased preference in positive conditions (z = 2.678, p = .007).

Finally, we ran a chi-square test of independence to test whether there was an association between people's future-biased and near-biased preferences. This test revealed that there was a significant association, χ^2 (4, N = 241) = 162.245, p < .001. However, this significant association disappears if we exclude people who report having a time-neutral preference, χ^2 (1, N = 172) = 1.781, p = .182. This suggests that the original association is being driven by the fact that people who reported having a time-neutral preference to future-biased prompts also tended to report having a time-neutral preference to near-biased prompts.[4]

2.2.3.2 *Experiment 2 Results.* We will begin by summarizing our major findings with respect to our hypotheses. Consider, first, the kind of sensation hypothesis according to which there will be differences in future bias (*H3*) and near bias (*H4*) across different kinds of sensation. Neither *H3* nor *H4* was supported. We found no evidence that future-biased or near-biased preferences were different across different kinds of sensation. Next, consider the negative

[4] Given the significant effect that valence exerts on people's future-biased and near-biased preferences, you might think that any potential association between these preferences might itself be associated with valence. To explore this possibility, we ran a Breslow-Day test. We found that there was a significant difference in the association between future-biased and near-biased preferences across valences, χ^2 (1, N = 241) = 9.577, p = .002. Importantly, the result persisted even when we excluded people who reported having time-neutral preferences, χ^2 (1, N = 172) = 5.614, p = .018. Future-biased and near-biased preferences are associated in positive conditions (p < .011), but not in negative conditions (p > .650).

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

valence hypothesis according to which people will be more future biased (*H6*) and more near biased (*H7*) about negative events than positive ones. Unlike in Experiment 1, we found no evidence in support of *H6* that more people have future-biased preferences regarding negative events than regarding positive ones. In contrast, just like in Experiment 1, *H7* was again not supported, with more people having far-biased preferences rather than near-biased preferences regarding negative events. Finally, consider the association hypothesis (*H5*) according to which there will be an association between future-biased preferences and near-biased preferences. Once again, we found no robust evidence of such an association.

Table 3 presents descriptive data of participants' responses regarding their future-biased preferences across all conditions in Experiment 2. The 'FB' column represents the number of participants who reported future-biased preferences. The 'PB' column represents the number of participants who reported past-biased preferences. The 'NP' column represents the number of participants who reported time-neutral preferences.

To investigate whether there was any difference in future-biased preferences across the conditions in Experiment 2, we ran a chi-square test of homogeneity. This test found no evidence that people's future-biased preferences differed across our conditions, χ^2 (14, N = 503) = 20.784, p = .107.

Table 4 summarizes the descriptive data of participants' responses regarding their near-biased preferences across all conditions. The 'NB' column represents the number of participants who reported near-biased preferences. The 'FrB' column represents the number of participants who reported far-biased preferences. The 'NP' column represents the number of participants who reported a time-neutral preference.

To investigate whether there was any difference of near-biased preferences across the conditions, we ran a chi-square test of homogeneity. The test showed that there was a significant difference in near-biased preferences across our conditions, χ^2 (14, N = 503) = 67.823, p < .001.

To identify the source of this effect we ran separate chi-square tests of homogeneity for valence and kind. The results of these tests found a significant effect of valence, χ^2 (2, N = 503) = 55.956, p < .001, and no significant effect of kind, χ^2 (6, N = 503) = 4.135, p = .658. Post-hoc comparisons with a Bonferroni

Table 3. Descriptive data of participants' responses to the future-biased prompt in Experiment 2

Valence	Kind	FB	PB	NP
Positive	Taste	19 (30.7%)	22 (35.5%)	21 (33.9%)
	Smell	22 (43.1%)	15 (29.4%)	14 (27.5%)
	Touch	36 (50.0%)	19 (26.4%)	17 (23.6%)
	Vision	24 (38.1%)	23 (36.5%)	16 (25.4%)
Negative	Taste	34 (56.7%)	13 (21.7%)	13 (21.7%)
	Smell	35 (51.5%)	22 (32.4%)	11 (16.2%)
	Touch	35 (53.9%)	22 (33.9%)	8 (12.3%)
	Vision	30 (48.4%)	22 (35.5%)	10 (16.1%)

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Table 4. Descriptive data of participants' responses to the near-biased prompt in Experiment 2

Valence	Kind	NB	FrB	NP
Positive	Taste	24 (38.7%)	22 (35.5%)	16 (25.8%)
	Smell	31 (60.8%)	8 (15.7%)	12 (23.5%)
	Touch	38 (52.8%)	21 (29.2%)	13 (18.1%)
	Vision	31 (49.2%)	15 (23.8%)	17 (27.0%)
Negative	Taste	13 (21.7%)	36 (60.0%)	11 (18.3%)
	Smell	12 (17.7%)	44 (64.7%)	12 (17.7%)
	Touch	20 (30.8%)	33 (50.8%)	12 (18.5%)
	Vision	16 (25.8%)	37 (59.7%)	9 (14.5%)

correction showed that people differed in their near-biased preferences in both negative and positive conditions. Standardized residuals showed that *more* people reported having a far-biased preference ($z = 8.635$, $p < .001$), and *fewer* people reported having a near-biased preference ($z = -3.188$, $p = .001$) or a time-neutral preference ($z = -5.447$, $p < .001$) in negative conditions. In contrast, *fewer* people reported having a time-neutral preference ($z = -3.323$, $p = .001$), and *more* people reported having a near-biased preference ($z = 5.568$, $p < .001$) in positive conditions.

Finally, we ran a chi-square test of independence to test whether there was an association between people's future-biased and near-biased preferences. This test revealed that there was a significant association, χ^2 (4, $N = 503$) = 190.886, $p < .001$. However, once again, this significant association disappeared when we excluded people who report having a time-neutral preference, χ^2 (1, $N = 364$) = 1.459, $p = .227$. This suggests that the original association is being driven by the fact that people who reported having a time-neutral preference to future-biased prompts also tended to report having a time-neutral preference to near-biased prompts.[5]

3. Discussion

3.1 Our Results and Preference-pattern Arguments

We had three broad hypotheses. First, we hypothesized that people's future-biased and near-biased preferences would be sensitive to the event factor. In particular, we tested the kind of hedonic event hypothesis according to which both near-biased and future-biased preferences are sensitive to whether the hedonic event in question is a sensation or a mood. And we tested the kind of sensation hypothesis according to which both near-biased and future-biased preferences are sensitive to the kind of sensation. Neither hypothesis was supported by our data.

[5] We ran a Breslow-Day test. We found no evidence that the association between future-biased and near-biased preferences differed across valences, χ^2 (1, $N = 503$) = 2.087, $p = .149$. This result does not change when we exclude people who report having time-neutral preferences, χ^2 (1, $N = 364$) = 0.01, $p = .919$.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Our second broad hypothesis was that people's near-biased and future-biased preferences will be sensitive to the valence factor. In particular, we tested the negative valence hypothesis according to which people will be more future- biased and more near- biased with regard to negative events than positive ones. That hypothesis was vindicated in the case of future bias (*H6*) where, in line with previous findings regarding future bias (Greene et al. 2021a, 2022a), across both experiments we found that future-biased preferences were stronger in the negative valence condition than in the positive one. We found the converse when it came to near-biased preferences, where it was *far*-biased preferences that were stronger in the negative valence condition. This, too, partially replicates previous findings of Latham, Miller, and Norton (2023) who found that people were more near biased and less far biased in positive conditions than in negative ones.

Our third broad hypothesis was that we would find an association between future-biased and near-biased preferences. This hypothesis was not supported. We failed to replicate the earlier findings of Latham, Miller, and Norton (2023) who found a moderate association between future-biased and near-biased preferences in both positive and negative conditions.

These results have implications for two current debates. But before we turn to these, we first want to draw attention to a limitation of the studies we ran. Our vignettes describe positive and negative experiences (sensations and moods in Experiment 1 and four kinds of sensations in Experiment 2) at the most general level without specifying the representational content or the specific qualia of these experiences. For instance, the participants were only told that they would experience a 'persistent, strong, extremely unpleasant visual effect'. This was deliberate because there may be variation in which particular visual effects are perceived as pleasant or unpleasant, and we wanted to eliminate misleading data resulting from individual idiosyncrasies. But the generality also invites a worry, namely, that when filling in the details of the (un)pleasant experiences, participants might have done so differently across the different valences.

For instance, perhaps participants were inclined to imagine negative experiences to be of greater *absolute* value than positive events even when the descriptions are symmetric ('persistent, strong, extremely (un)pleasant'). For instance, some participants might have imagined the positive visual experience to be one that lacks any visual anomalies (such as blurriness, short-sightedness, flashes, floaters, and so on) and one that has additional clarity and perhaps richness of color, but imagined the negative experience to be one that contains all of the visual anomalies just listed and then some. In other words, while the 'extremely pleasant' visual experience might have been perceived as somewhat better than neutral, the 'extremely unpleasant' experience might have been perceived as much worse than neutral.

This potential confound could be obscuring the real role of valence. This possibility casts some doubt on our results regarding *H6*. However, while this hypothesis could explain why people were more future biased about negative events than about positive ones even if they are not in fact sensitive to valence, it

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

could not explain our results regarding *H7*. Indeed, if this were the explanation for our results regarding *H6*, we would expect *H7* to be vindicated: people would be more near biased about negative events than about positive ones. There should be no asymmetry between the perceived values of the experiences at the population level in conditions of near bias and of future bias. Accordingly, we set this concern aside.

Let us return to the implications of our finding for two current debates. First, our results do not support a general preference-pattern argument that appeals to the event factor. That argument would go as follows:

Generalized Event Factor Preference-Pattern Argument

13. Preferences that are sensitive to normatively irrelevant factors are objectionably arbitrary.
14. We have reason to avoid having objectionably arbitrary preferences.
15. Future-biased and near-biased preferences are both sensitive to the event factor.
16. The event factor is normatively irrelevant.
17. Therefore, future-biased and near-biased preferences are both sensitive a normatively irrelevant factor (from 15, 16).
18. Therefore, we have reason to avoid having future-biased or near-biased preferences (from 13, 14, 17).

We found no evidence in favor of (15). This does not definitively show that such preferences are *not* sensitive to the event factor. There could be another way in which these preferences are sensitive to types of event that our experiments did not test. Nonetheless, our results cast some doubt on the argument.

Next, let us turn to the argument as it applies to valence.

Generalized Valence Factor Preference-Pattern Argument

19. Preferences that are sensitive to normatively irrelevant factors are objectionably arbitrary.
20. We have reason to avoid having objectionably arbitrary preferences.
21. Future-biased and near-biased preferences are both sensitive to the valence factor.
22. The valence factor is normatively irrelevant.
23. Therefore, future-biased and near-biased preferences are both sensitive to a normatively irrelevant factor (from 21, 22).
24. Therefore, we have reason to avoid having future-biased or near-biased preferences (from 19, 20, 23).

If this argument succeeds, it puts some pressure on normative hybridism. Our results shed light on the status of this argument. We found that people's near-biased and future-biased preferences are sensitive to valence. That begins to

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

put this argument on a firmer footing. In order for the argument to succeed, however, advocates of time-neutrality need to argue that (22) is true: valence is normatively irrelevant.

We think this claim is *intuitive*. Consider the following pair of examples. In the first, you are asked for your preference between either having 20 units of disutility yesterday or 10 tomorrow. In the second, you are asked for your preference between either having 20 units of utility yesterday or 10 tomorrow. Now suppose further that you are future biased in both cases, but that you are more strongly future biased in the first case than in the second. It is very hard to see why the strength of your preference is not arbitrary in that case. Why should the fact that it is disutility in one case and utility in the other make any difference here?

Of course, it is not so easy to be sure that people are sensitive to valence in this manner. It is consistent with our results that three hours of pain *seem more bad* to many people than three hours of pleasure seem good. Thus, while people's preference might be sensitive to valence, this could be because negatively valenced events are (or are regarded as being) more disutilitous than positively valenced ones are utilitous. Nonetheless, there seems to be significant scope for time-neutralists to develop this valence preference-pattern argument in order to argue that we have reason to avoid both near bias and future bias.

3. 2 Some Complications: Anticipation and Retrospection

One curious feature of the valence-sensitivity of near bias and future bias, however, is that it pulls in opposite directions. Consistent with Greene et al. (2021a, 2022a), we found stronger future-biased preferences in negative conditions than in positive conditions; however, contrary to our prediction (*H7*), we found stronger near bias in *positive* conditions (see also Molouki, Hardisty, and Caruso 2019; Hardisty and Weber 2020). This might be seen as undermining the generalized valence factor preference-pattern argument by cutting against the assumption that valence is normatively irrelevant (22). If valence drives our preferences through some general affective mechanism that we might think is normatively suspicious by, say, undermining cool rationality, then one would expect valence to have a uniform effect on near-biased and future-biased preferences. That valence does not have a uniform effect suggests that there is no single mechanism and at least opens up the possibility that valence might be normatively relevant.

This thought connects to another finding of our study—the absence of an association between near-biased and future-biased preferences (*contra* Latham, Miller, and Norton 2023)—and to the second debate with respect to which our results are relevant, namely, the debate concerning the explanation for future-biased and near-biased preferences. A lack of association between future bias and near bias would speak against a shared explanation and against tying together the normative status of these biases. Because we found opposite effects of valence, this might be seen as support for normative hybridism. However, in order to interpret these findings, we may have to move beyond the surface-level patterns

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

of future bias and near bias (e.g., the [lack of] association and the factors they are sensitive or insensitive to) and further investigate the underlying mechanisms. It could be that there is no association, as our results suggest. But it could also be that methodological differences in the relevant studies are obscuring things. We will now discuss one significant potential source of obscurity.

Studies in this area have sought to control for various factors, including the probability of the relevant events, their subjective value to the person who experiences them, and so on. But no studies have aimed to diminish negative or positive anticipation or retrospection where by 'anticipation' we mean the conscious state of imagining and entertaining oneself undergoing certain future events, and by 'retrospection' we mean the conscious state of imagining and entertaining oneself having undergone certain past events. By extension, we will say that positive anticipation is the anticipation of a positive event, and negative anticipation is the anticipation of a negative event and *mutatis mutandis* for positive/negative retrospection.

To be sure, many of these studies appeal to amnesia (though the study by Lee et al. [2020] does not mention amnesia). In Parfit's original thought experiment you wake up in hospital and cannot remember whether you just had a painful operation or are still to have it. In the experiment by Greene et al. (2020) you wake up and *for a moment* cannot remember whether you already ate your favorite/most disliked meal. In most of these experiments while there is no *retrospection* of the event in question at the time the preference is being adduced, there is likely anticipation of the event (if it will be in the future) and there may also be anticipation of retrospection at a later time.

Even if we can construct realistic cases in which anticipation and retrospection are stipulated to be absent, it might still be unavoidable that anticipation and retrospection influence our preferences. To take an example, in Parfit's presentation of the case about the past and future operations (1984: 165–66), he stipulated that the memory of the past operation would be removed such that retrospection is absent, but there is no parallel stipulation that the reader would not anticipate the *future* operation. To 'make all else equal,' Brink (2011) suggests one way to modify the case: to 'change the example so that it involves administration of a drug that blocks anticipation of future pain, much as the doctors induce amnesia to block recollection of the pain of the operation' (2011: 379).

There are two reasons to suspect that this modification will not work as intended. First, it seems impossible to form a temporal preference involving future events without anticipating the future event at all. In this example, you must learn the alternatives and entertain how good or bad they are for you, without which there would be no preference; therefore, some anticipation must be present on pain of rendering the example unintelligible. (Likewise, it might be suggested that you must imagine having a painful memory in order to evaluate whether a past surgery is preferable. (cf. Phillips 2021)) Second, there is a subtle difference between forming a preference *from* that scenario and forming a surrogate preference *regarding an imagined* scenario. Although the imagined 'you' is supposed to lack anticipation, the 'you' who is reading the vignette would still be

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

inclined to anticipate the future operation because this is what we ordinarily do. In other words, surrogate anticipation would contaminate the formation of your preference. (A similar case can be made for surrogate retrospection.)

In light of this, one may suggest that we instead make anticipation and retrospection equally present and let them 'cancel out' each other. Brink considers this alternative: 'We must change the example so that the past suffering is something that one can recollect, just as prospective pain can be anticipated' (2011: 379). While this way of making 'all else equal' sounds more promising, it is difficult to make anticipation and retrospection of equal emotional intensity because (1) people tend to experience stronger emotions when anticipating (we shall elaborate on this point shortly), and (2) emotional intensity also depends on the temporal distances of the events and many other factors.

For all we know, in the absence of valenced anticipatory or retrospective states, people would not exhibit future bias. It is not easy to see how one would empirically test this. Although it is relatively easy to describe scenarios in which there is no retrospection, it is less clear how to describe scenarios in which there is no positive or negative anticipation of events that are positively or negatively valenced, and it is not clear whether people can even form preferences for the temporal locations of events that they do not anticipate.

There are really two concerns here. One is a concern about whether we could ever *isolate* and *test for* genuine near bias and future bias (which are only sensitive to the temporal locations of the alternatives). If we cannot isolate genuine near bias or future bias, then we cannot even be sure that people exhibit such biases. Another concern, though, is that perhaps talk of genuine future bias or near bias makes little sense. Callender (2021) articulates this kind of worry when it comes to near bias. He notes that a whole range of properties—from aging and memory to thermodynamic properties—is connected to temporal locations. He notes how difficult it is for experiments to determine whether it is these factors or temporal location itself driving peoples' preferences. In addition, he is not sure whether it makes sense to think of our preferences as sensitive to temporal location itself, rather than to these properties of temporal locations. The worry is that there is no sense to be made of our having such preferences 'holding all else equal' because it is not even clear that we can have preferences at all if we hold all else equal. And the same is true of anticipation. It is not clear what sorts of creatures we would need to be in order for our anticipatory states to play no role in guiding our preferences. Thus, it is not clear whether sense can be made of genuine near bias or future bias understood as a preference we have holding all else equal, where the 'all else' here includes retrospective and anticipatory states.

We take no stance on whether the 'all else being equal' should include facts about anticipation and retrospection. But whichever stance one takes, interesting questions remain regarding the normative status of the resulting preferences and, in particular, regarding whether positive and negative anticipation/retrospection give us normative reason to have *apparently* time-biased preferences (i.e., preferences that are either merely apparently or genuinely time-biased).

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

3.3 A Hypothesis about Anticipation and Retrospection and their Connection to Additional Valenced Mental States

We hypothesize that the effects of positive/negative anticipation/retrospection are what partially explain our failure to vindicate (1) the association hypothesis vindicated by Latham, Miller, and Norton (2023) and (2) the negative valence hypothesis in apparent near biases. States of anticipation/retrospection tend to be valenced. For instance, the experience of anticipating a negative event may be unpleasant; *mutatis mutandis* for a positive event. It could also be that states of anticipation or retrospection generate, or are associated with, further valenced mental states. For instance, anticipating a negative event may generate anxiety. Imagine Freddie, who is sitting in the dentist's waiting room waiting for a procedure. Freddie feels anxiety: he paces the room, his heart rate quickens, his palms are sweaty, his stomach is upset. The retrospection of negative events may similarly generate additional valenced mental states (such as fear, mortification, anxiety, and so on). As with anticipation, research shows that recollection can itself be more or less pleasurable (Elster and Loewenstein 1992; Morewedge 2015). However, the net utility of anticipation or retrospection need not be congruent with that of the objects of anticipation or retrospection for there are also ways in which *positive/negative* anticipation/retrospection can give rise to *negative/positive* present emotions.

The consumption effect occurs when pleasant experiences are derived from positive anticipation/retrospection or unpleasant experiences are derived from negative anticipation/retrospection (Elster and Loewenstein 1992: 216–17; Loewenstein 1987). When anticipating a positive or negative event, especially when the imagination is vivid and emotionally intense, pleasant or unpleasant feelings tend to be present as if the event were experienced in advance. When recollecting or imagining having undergone a positive or negative event, likewise, the states of recollection and imagination can also be pleasant or unpleasant.

The contrast effect (Elster and Loewenstein 1992: 216–17) occurs when a comparison between a current state and some other state (your own or someone else's) leads to positive/negative experiences, such as when your pleasant anticipation of a trip is reduced upon your comparing it with a fancier trip that you are not taking.

Suppose that only the consumption effect is present in positive/negative anticipation/retrospection. Then the anticipation of positive/negative events is in itself pleasant/unpleasant. In such a case the consumption effect *decreases* apparent (prospective) near-biased preferences. We will call a reason for preferences that is generated in this way a *simple motivational reason*.

There is another way in which the consumption effect decreases apparent near-biased preferences. When forming a temporal preference, one may take into consideration the expected utility of anticipation—that is, savoring and dread (in the case of positive and negative events, respectively)— in addition to the (subjective) utilities of the future events in question (Loewenstein 1987; Lee et al. 2022). If one is motivated to have certain temporal preferences in part due to considerations arising from the consumption effect, one has a *considered motivational reason*.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

What about positive/negative retrospection? First consider the consumption effect of retrospection taken in isolation. Again, there are two ways in which the consumption effect of retrospection leads to *decreases* in apparent future-biased preferences. When one recalls a positive/negative event in the past or imagines that a positive/negative event has happened, the retrospection is pleasant/unpleasant. Consider a pleasant retrospection. One might react to the pleasant retrospection by desiring that it be present rather than absent. Alternatively, one might factor in the utility/disutility of pleasant/unpleasant retrospections when one determines the overall utilities of the alternatives in forming the temporal preference. Regardless of whether one is motivated by the simple reason or the considered reason, then, the consumption effect of retrospection leads to *decreases* in apparent future biases.

Because the consumption effect of anticipation results in *increases* in apparent future-biased preferences for parallel reasons, it might be expected that the motivational forces of anticipation and retrospection simply cancel out each other and result in past-/future-neutral preferences. We know, however, that all else being equal, people tend to experience more intense emotions during anticipation than during retrospection of the same experience (Caruso, Gilbert, and Wilson 2008; D'Argembeau and Van der Linden 2004; Van Boven and Ashworth 2007), and our emotional reactions to consideration of past experiences are less extreme than our reactions to consideration of future experiences (Van Boven, White, and Huber 2009). Given this, we should predict that, on balance, the preference for the positive/negative future event to be present/absent rather than absent/present would outweigh the preference for the positive/negative past event to be present/absent rather than absent/present. And we should also expect that in cases in which we find higher levels of positive and negative anticipation and retrospection, we should find *increases* in apparent future biases due to greater unbalance of emotional intensity.

The contrast effect, on the other hand, results in *increases* in apparent (prospective) near biases. Suppose that only the contrast effect is operative in anticipation of future events. When one anticipates a pleasant/unpleasant future event, the utility/disutility in the present is derived from the states of anticipation because the cross-temporal comparison leads one's present condition to be perceived as more/less satisfactory. A notable feature of the contrast effect in anticipation is that although it is well documented in positive anticipation, it is rarely observed in negative anticipation (Hardisty and Weber 2020: 610; Molouki, Hardisty, and Caruso 2019: 1676; Liberman et al. 2009; Affleck et al. 2000; Brickman, Coates, and Janoff-Bulman 1978).

Taking all these factors into account, anticipation tends to result in a mixture of increases and decreases in apparent near biases for positive events due to both effects working in tandem, whereas in negative conditions, because the consumption effect dominates the contrast effect, anticipation almost uniformly leads to decreases in apparent near biases. This asymmetry between positive and negative anticipation, we hypothesize, is what explains the following results in our experiments. First, the negative valence hypothesis was not vindicated in the case of apparently near-biased preferences. Second, the converse was true: we found more apparently

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

far-biased preferences with regard to negative events. Our suggestion is similar to the explanation for the sign effect proposed by Molouki, Hardisty, and Caruso (2019) and Hardisty and Weber (2020), where the sign effect refers to the phenomenon that 'people discount future positives more than future negatives' (Molouki, Hardisty, and Caruso 2019: 1674); the sign effect is exactly what we have found in our experiments described above.

We have suggested that when it comes to negative events the consumption effect dominates, whereas the consumption effect and the contrast effect are both efficacious with regard to positive events. This means that when positive and negative anticipation are strong, we should expect *a larger* decrease in apparently near-biased preferences for negative events. Because we do not know the relative weights of the consumption and contrast effects in the positive condition, all we can say is that holding the increased strength of anticipation fixed, the variation in (either increases or decreases) apparent near biases for positive events would be significantly less than the decreases in apparent near biases for negative events. Even so, this still leaves it mysterious that when anticipation is plausibly quite weak (as in Latham, Miller, and Norton 2023) apparent near-biased preference is *stronger* in negative than in positive conditions.

Here we will offer two alternative hypotheses. First, it may be that other factors are at play: perhaps in general we discount future negatives more than future positives because, for example, we feel temporally more distant from future selves who are worse off than those who are better off, such that when anticipation is weak, the negative valence hypothesis would be vindicated. But when anticipation is strong, the contrast effect in positive anticipation tips the scale. Or, alternatively, perhaps there are some differences between their vignettes and ours that have driven people to be more apparently near biased in their study but more apparently far biased in our in negative conditions. We are not sure what those factors might be, and further work could profitably be directed at this question.

To sum up then, we have suggested that the consumption effect in anticipation/retrospection leads to increases/decreases in apparently future-biased preferences, and because anticipation tends to be emotionally more intense than retrospection, on balance, the consumption effect tends to result in increases in apparently future-biased preferences. We have also seen that the contrast effect is salient in positive anticipation but rarely observed in negative anticipation. Thus, taken in isolation the contrast effect leads to *increases* in apparent near biases. For similar reasons, the contrast effect leads to *decreases* in apparent future biases. The contrast effect in positive and negative retrospection, on the other hand, leads to *increases* in apparent future biases. However, the contrast effects in retrospection and anticipation do not simply cancel out each other in both negative and positive conditions. While the contrast effect in retrospection for both positive events and negative events is found to be a usual occurrence, the contrast effect is barely existent in negative anticipation. On balance, we should expect that the contrast effect, taken in isolation, leads to substantial increases in apparently future-biased preferences for *negative* events, whereas its effect on future-biased preferences for positive events should be relatively small. And this,

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

we hypothesize, is what explains the vindication of the negative valence hypothesis in apparently future-biased preferences.

All this suggests that further work could profitably be undertaken separating out these effects to determine their relative importance and perhaps also to compare those effects in different conditions (such as first-person versus third person and hedonic versus non-hedonic conditions). It also suggests that further work could be directed at thinking about the normative role of these two kinds of effects.

Finally, we should consider whether there is an asymmetry between anticipation and retrospection when it comes to grounding the rational permissibility of time biases. Interestingly, many philosophers seem to have thought that holding all else equal the fact that we do *not* anticipate past events is a reason to discount their value relative to future events despite the fact that we *do* retrospect those events. These philosophers also think that the fact that we less strongly anticipate events that are farther in the future, compared to those that are nearer, is *not* reason to discount the events' value relative to the near. But it is not obvious what could ground this asymmetry. Perhaps some of these philosophers have thought that we *should not* anticipate past events (and we do not) and that this is why future bias is rationally permissible (or obligatory). But it is also the case that we more strongly anticipate temporally near events over temporally more distant events. Accordingly, it seems that those who defend normative hybridism would either need to argue that (a) anticipation does not play a role in grounding the normative status of near bias although it does play such a role in grounding the normative status of future bias or (b) that we *should equally strongly* anticipate far future events and nearer ones or (c) that reasons arising from the nature of anticipation afford near bias and future bias the same normative status, but that in the case of near bias there are other countervailing reasons showing that near bias is in fact rationally impermissible. For instance, defenders of normative hybridism might take option (c) and argue that, as Moller (2002: 77) points out, one key difference between the two biases is that while it is obvious that we can trade between the near and distant future, it is not obvious that we can trade between the future and the past. Thus, if the rational status of these biases is tied to the extent to which having these preferences can affect our choices or their outcomes, this may be reason to think that near bias is irrational and future bias is not.

3.4 Limitations of our Studies

Before concluding, it is worth considering some limitations of these studies. Some readers might worry that the vignettes and questions were too cognitively demanding for nonphilosophers to understand, and so our results might not indicate much about people's temporal preferences. To address this concern, we included both attention checks and comprehension check questions. These questions served two purposes. First, they weeded out bots and people selecting answers at random or without thought in order to receive payment quickly, something that needs to be guarded against when running studies online (Ahler,

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Roush, and Sood 2020). Second, they enabled us to achieve a sample composed of people who understood both the vignettes and the questions being asked.

Some readers might then worry that even if our current results are informative, they might only be narrowly so because the remaining sample is not representative. After all, a little over the half the sample had to be excluded from the analyses in each experiment. It is worth noting that high exclusions are not uncommon in experimental philosophy studies that include comprehension check questions. For example, Nichols and Bruno (2010) found that 33 percent and 44 percent of students in a University of Arizona undergraduate philosophy class failed to understand a personal identity thought experiment adequately. More recently, Nadelhoffer, Murray, and Murry (2021) investigated subjects' understanding of determinism and found that 81 percent failed to understand it correctly. *Perhaps* those who pass our comprehension checks are more reflective, thoughtful, and so on, than those who do not. Even if that is right, we have no reason at present to think that such people have different temporal preferences than others.

Finally, some readers might worry that our current results are only narrowly informative because the Mechanical Turk population is not representative of the broader population. There is evidence that Turkers are younger, more educated, have lower incomes, and are less racially and ethnically diverse than the US average (e.g., Difallah, Filatova, and Ipeirotis 2018; Hitlin 2016). Nevertheless, the Mechanical Turk population is more representative than typical samples of convenience (i.e., undergraduate university students). Still, some caution is warranted regarding the generalizability of these findings. For example, there is evidence that stronger near-biased preferences are associated with lower wealth and higher income inequality (e.g., Epper et al. 2020; Ludwig, Flournoy, and Berkman 2019; Ruggeri et al. 2022). Unfortunately, possible differences along demographic lines are not something that we controlled for in this study. Thus, it would be good to rerun this study with a more representative sample and target demographic characteristics that might be associated with differences in temporal preferences. For instance, it would be interesting to examine whether future-biased preferences are associated with the same demographic factors as near-biased preferences and in the same fashion. It would also be interesting to investigate whether the association between future-biased preferences and near-biased preferences is itself associated with demographic factors.

4. Conclusion

We think that future work in this area could profitably be directed at answering four questions. First, are near bias and future bias genuine only if 'all else being equal' includes facts about anticipation and retrospection? Second, what role do anticipation and retrospection play in our having near-biased and future-biased preferences? Third, what normative status do these states of anticipation and retrospection tend to confer on these preferences? And, fourth, is there some asymmetry in the normative status of anticipation and retrospection, an

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

asymmetry which might explain the rational permissibility of future bias but not near bias, i.e., normative hybridism?

ANDREW J. LATHAM
AARHUS UNIVERSITY, AARHUS, DENMARK
andrew.latham@cas.au.dk

KRISTIE MILLER
UNIVERSITY OF SYDNEY, SYDNEY, NSW, AUSTRALIA
Kristie.miller@sydney.edu.au

JORDAN OH
UNIVERSITY OF SYDNEY, SYDNEY, NSW, AUSTRALIA
Veoh7380@uni.sydney.edu.au

SAM SHPALL
UNIVERSITY OF SYDNEY, SYDNEY, NSW, AUSTRALIA
sam.shpall@sydney.edu.au

WEN YU
UNIVERSITY OF SYDNEY, SYDNEY, NSW, AUSTRALIA
weyu9217@uni.sydney.edu.au

References

Affleck, Glenn, Howard Tennen, Susan Urrows, Pamela Higgins, and Micha Abeles. (2000) 'Downward Comparisons in Daily Life with Chronic Pain: Dynamic Relations with Pain Intensity and Mood'. *Journal of Social and Clinical Psychology,* 19, 499–518. https://doi.org/10.1521/jscp.2000.19.4.499.

Ahler, Douglas J., Carolyn E. Roush, and Gaurav Sood. (2021) 'The Micro-Task Market for Lemons: Data Quality on Amazon's Mechanical Turk'. *Political Science Research and Methods*, October, 1–20. https://doi.org/10.1017/psrm.2021.57.

Ainslie, G., and Haslam, N. (1992) 'Self-control'. *Choice Over Time,* 177, 209.

Bickel, Warren K., Richard Yi, Benjamin P. Kowal, and Kirstin M. Gatchalian. (2008) 'Cigarette Smokers Discount Past and Future Rewards Symmetrically and More than Controls: Is Discounting a Measure of Impulsivity?' *Drug and Alcohol Dependence,* 96, 256–62. https://doi.org/10.1016/j.drugalcdep.2008.03.009.

Bramble, Ben. (2016) 'A New Defense of Hedonism about Well-Being'. *Ergo, an Open Access Journal of Philosophy*, 3. https://doi.org/10.3998/ergo.12405314.0003.004.

Brickman, Philip, Dan Coates, and Ronnie Janoff-Bulman. (1978) 'Lottery Winners and Accident Victims: Is Happiness Relative?' *Journal of Personality and Social Psychology,* 36, 917–27. https://doi.org/10.1037/0022-3514.36.8.917.

Brink, David O. (2011). 'Prospects for Temporal Neutrality'. In Craig Callender (ed.), *The Oxford Handbook of Philosophy of Time* (Oxford: Oxford University Press), 353–81.

Callender, Craig. (2021) 'The Normative Standard for Future Discounting'. *Australasian Philosophical Review,* 5, 227–53. https://doi.org/10.1080/24740500.2021.2112087.

Caruso, Eugene M., Daniel T. Gilbert, and Timothy D. Wilson. (2008) 'A Wrinkle in Time: Asymmetric Valuation of Past and Future Events'. *Psychological Science,* 19, 796–801. https://doi.org/10.1111/j.1467-9280.2008.02159.x.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Craig, William Lane. (1999) 'Tensed Time and Our Differential Experience of the Past and Future'. *Southern Journal of Philosophy,* 37, 515–37. https://doi.org/10.1111/j.2041-6962.1999.tb00880.x.

D'Argembeau, Arnaud. and van der Linden, Martial. (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, 844–58.

Difallah, Djellel, Elena Filatova, and Panos Ipeirotis. (2018) 'Demographics and Dynamics of Mechanical Turk Workers'. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 135–43. https://doi.org/10.1145/3159652.3159661.

Dougherty, Tom. (2015) 'Future-Bias and Practical Reason'. *Philosophers' Imprint,* 15, 1–16.

Elster, John, and George Loewenstein. (1992) 'Utility from Memory and Anticipation'. In G. Loewenstein and J. Elster (eds.), *Choice over Time* (New York: Russell Sage Foundation), 213–34.

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Søren Leth-Petersen, and Gregers Nytoft Rasmussen. (2020) 'Time Discounting and Wealth Inequality'. *American Economic Review,* 110, 1177–1205. https://doi.org/10.1257/aer.20181096.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. (2002) 'Time Discounting and Time Preference: A Critical Review'. *Journal of Economic Literature*, 40, 351–401. https://doi.org/10.1257/jel.40.2.351.

Greene, Preston, Andrew J. Latham, Kristie Miller, and James Norton. (2020) 'Hedonic and Non-Hedonic Bias Toward the Future'. *Australasian Journal of Philosophy,* 99, 1–16. https://doi.org/10.1080/00048402.2019.1703017.

Greene, Preston, Andrew J. Latham, Kristie Miller, and James Norton. (2021a) 'On Preferring That Overall, Things Are Worse: Future-Bias and Unequal Payoffs'. *Philosophy and Phenomenological Research,* 105, 181–94. https://doi.org/10.1111/phpr.12819.

Greene, Preston, Andrew J. Latham, Kristie Miller, and James Norton. (2021b) 'The Rationality of Near Bias Toward Both Future and Past Events'. *Review of Philosophy and Psychology,* 12, 905–22. https://doi.org/10.1007/s13164-020-00518-1.

Greene, Preston, Andrew J. Latham, Kristie Miller, and James Norton. (2022a). 'Capacity for Simulation and Mitigation Drives Hedonic and Non-Hedonic Time Biases'. *Philosophical Psychology*, 35, 226–52. https://doi.org/10.1080/09515089.2021.1960299.

Greene, Preston, Andrew J. Latham, Kristie Miller, and James Norton. (2022b) 'How Much Do We Discount Past Pleasures?' *American Philosophical Quarterly,* 59, 367–76. https://doi.org/10.5406/21521123.59.4.03.

Greene, Preston, Andrew J. Latham, Kristie Miller, James Norton, Christian Tarsney, and Hannah Tierney. (2022c) 'Bias Towards the Future'. *Philosophy Compass,* 17, e12859. https://doi.org/10.1111/phc3.12859.

Hardisty, David J., and Elke U. Weber. (2020) 'Impatience and Savoring vs. Dread: Asymmetries in Anticipation Explain Consumer Time Preferences for Positive vs. Negative Events'. *Journal of Consumer Psychology,* 30, 598–613. https://doi.org/10.1002/jcpy.1169.

Hare, Caspar. (2013) 'Time – The Emotional Asymmetry'. In Heather Dyke and Adrian Bardon (eds.), *A Companion to the Philosophy of Time* (Chichester: John Wiley), 507-20. https://doi.org/10.1002/9781118522097.ch29.

Hausman, J. A. (1979) 'Individual Discount Rates and the Purchase and Utilization of Energy-using Durables'. *Bell Journal of Economics,* 10, 33–54.

Heathwood, Chris. (2008) 'Fitting Attitudes and Welfare'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics 3* (Oxford: Oxford University Press), 47–73.

Hedden, Brian. (2015). *Reasons Without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.

Hitlin, Paul. (2016) 'Research in the Crowdsourcing Age, a Case Study'. Pew Research Center. http://www.pewinternet.org/2016/07/11/research-in-the-crowdsourcing-age-a-case-study/.

Kauppinen, Antti. (2018) 'Agency, Experience, and Future Bias'. *Thought: A Journal of Philosophy*, 7, 237–45.

https://doi.org/10.1017/apa.2023.14 Published online by Cambridge University Press

Labukt, Ivar. (2012) 'Hedonic Tone and the Heterogeneity of Pleasure'. *Utilitas,* 24, 172–99. https://doi.org/10.1017/S0953820812000052

Latham, Andrew J., Kristie Miller, and James Norton. (2023) 'Against a Normative Asymmetry between Near- and Future-Bias'. *Synthese,* 201, 93. https://doi.org/10.1007/s11229-023-04045-1.

Lee, Ruth, Christoph Hoerl, Patrick Burns, Alison S. Fernandes, Patrick A. O'Connor, and Teresa McCormack. (2020) 'Pain in the Past and Pleasure in the Future: The Development of Past–Future Preferences for Hedonic Goods'. *Cognitive Science,* 44, e12887. https://doi.org/10.1111/cogs.12887.

Lee, Ruth, Jack Shardlow, Patrick A. O'Connor, Lesley Hotson, Rebecca Hotson, Christoph Hoerl, and Teresa McCormack. (2022) 'Past-Future Preferences for Hedonic Goods and the Utility of Experiential Memories'. *Philosophical Psychology,* 35, 1181–1211. https://doi.org/10.1080/09515089.2022.2038784.

Liberman, Varda, Julia K. Boehm, Sonja Lyubomirsky, and Lee D. Ross. (2009) 'Happiness and Memory: Affective Significance of Endowment and Contrast'. *Emotion,* 9, 666–80. https://doi.org/10.1037/a0016816.

Loewenstein, George. (1987) 'Anticipation and the Valuation of Delayed Consumption'. *The Economic Journal,* 97, 666. https://doi.org/10.2307/2232929.

Lowry, Rosemary, and Martin Peterson. (2011) 'Pure Time Preference'. *Pacific Philosophical Quarterly,* 92, 490–508. https://doi.org/10.1111/j.1468-0114.2011.01408.x.

Ludwig, Rita M., John C. Flournoy, and Elliot T. Berkman. (2019) 'Inequality in Personality and Temporal Discounting across Socioeconomic Status? Assessing the Evidence'. *Journal of Research in Personality,* 81, 79–87. https://doi.org/10.1016/j.jrp.2019.05.003.

Moller, Dan. (2002) 'Parfit on Pains, Pleasures, and the Time of Their Occurrence'. *Canadian Journal of Philosophy,* 32, 67–82.

Molouki, Sarah, David J. Hardisty, and Eugene M. Caruso. (2019) 'The Sign Effect in Past and Future Discounting'. *Psychological Science,* 30, 1674–95. https://doi.org/10.1177/0956797619876982.

Morewedge, Carey. (2015) 'Utility: Anticipated, Experienced, and Remembered'.' In G. Keren and G. Wu (eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making* (Hoboken, NJ: Wiley & Sons), 295–330. https://doi.org/10.1002/9781118468333.ch10.

Nadelhoffer, Thomas, Samuel Murray, and Elise Murry. (2021) 'Intuitions about Free Will and the Failure to Comprehend Determinism'. *Erkenntnis,* 88, 2515–36. https://doi.org/10.1007/s10670-021-00465-y.

Nichols, Shaun, and Michael Bruno. (2010) 'Intuitions about Personal Identity: An Empirical Study'. *Philosophical Psychology,* 23, 293–312. https://doi.org/10.1080/09515089.2010.490939.

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

Pearson, Olley. (2018) 'Appropriate Emotions and the Metaphysics of Time'. *Philosophical Studies,* 175, 1945–61.

Phillips, Callie K. (2021) 'Why Future-bias isn't Rationally Evaluable'. *Res Philosophica,* 98, 573–96.

Prior, A. N. (1959) 'Thank Goodness That's Over'. *Philosophy,* 34, 12–17. doi:10.1017/S0031819100029685.

Rawls, John. (1971) *A Theory of Justice*. Cambridge, MA: Belknap.

Ruggeri, K., B. Većkalov, R. Freichel, S. J. Geiger, and E. García-Garzon. (2022) 'The Globalizability of Temporal Discounting'. *Nature Human Behavior,* 6, 1386–97.

Schlesinger, George. (1976) 'The Stillness of Time and Philosophical Equanimity'. *Philosophical Studies,* 30, 145–59.

Sidgwick, Henry. (1884) *The Methods of Ethics*. London: Macmillan

Soman, D., Ainslie, G., Frederick, S., Li, X., Lynch, J., Moreau, P., Mitchell, A., Read, D., Story, G. W., Vlaev, I., Seymour, B., Darzi, A., and Dolan, R. J. (2014) Does Temporal Discounting Explain Unhealthy Behavior? A Systematic Review and Reinforcement Learning Perspective. *Frontiers in Behavioral Neuroscience,* 8, 76. https://doi.org/10.3389/fnbeh.2014.00076

Thaler, R. (1981) 'Some Empirical Evidence on Dynamic Inconsistency'. *Economics Letters,* 8, 201–7.

Trout, J. D. (2007). 'The Psychology of Discounting: A Policy of Balancing Biases'. *Public Affairs Quarterly*, 21, 201–20.

Van Boven, L., and Ashworth, L. (2007) 'Looking Forward, Looking Back: Anticipation is More Evocative than Retrospection'. *Journal of Experimental Psychology: General*, 136, 289–300. https://doi.org/10.1037/0096-3445.136.2.289

Van Boven, L., White, K., and Huber, M. (2009) 'Immediacy Bias in Emotion Perception: Current Emotions Seem More Intense Than Previous Emotions'. *Journal of Experimental Psychology. General*, 138, 368–82. https://doi.org/10.1037/a0016074

Yi, Richard, Kirstin M. Gatchalian, and Warren K. Bickel. (2006) 'Discounting of Past Outcomes'. *Experimental and Clinical Psychopharmacology*, 14, 311–17. https://doi.org/10.1037/1064-1297.14.3.311.