

Discounting the past: past/future symmetries in temporal preferences?

Abstract

A person is *near biased* if, all else being equal, they prefer positive events to be near rather than distant and negative events to be distant rather than near. The most prominent form of the bias is *future-directed*. A future near-biased agent prefers positive events to be in their near future rather than their distant future and negative events to be in their distant future rather than their near future. To date, experimental work has focused on future-directed near bias. In this paper we empirically investigate *past-directed* near bias. As predicted, we found a population-level symmetry between past-directed and future-directed near bias across both hedonic and non-hedonic events. That is, we see the same pattern of preferences about future events as we see regarding past events. Further, we added to earlier results that show a difference in time bias between preferences for hedonic and non-hedonic events. Startlingly, however, despite the population-level pattern of past/future symmetry, we found no correlation between individual participants' past- and future-directed preferences. This suggests that the mechanisms that underlie the pattern of preferences regarding future events are different from the mechanisms that underlie the pattern of preferences regarding past events.

1. Introduction

A person is *near biased* when, all else being equal, they prefer positive events to be near rather than distant and negative events to be distant rather than near. The most prominent form of the bias is *future-directed*. A future near-biased agent prefers positive events to be in their near future rather than their distant future and negative events to be in their distant future rather than their near future.

To date, experimental work has focused on the descriptive project of investigating the patterns of future near bias people exhibit. Philosophers and social scientists have also been interested in the normative status of near-biased preferences. Social scientists study future near bias under the label *temporal discounting*, the idea being that the further into the future an event is, the more agents discount its value. Normatively, social scientists tend to claim that future near bias (i.e., temporal discounting) is rationally permissible only if the agent's preferences are consistent over time, while philosophers tend to claim that future near bias is irrational.¹

Another way a person can be near biased is *past-directed*. A past near-biased agent prefers positive events to be in their near past rather than their distant past and negative events to be in their distant past rather than their near past. There is very little experimental work focussing on past near bias, and philosophers have paid no attention to what little work there is. But any theory of the rational permissibility or impermissibility of near bias ought to have something to say about past near bias. As with future near bias, experimental research is beneficial in determining the normative

¹ See Greene and Sullivan (2015:§2). Sidgwick (1884:380–1) writes, “The mere difference of priority and posteriority in time is not a reasonable ground for having more regard to the consciousness of one moment than to that of another. The form in which it practically presents itself to most men is ‘that a smaller present good is not to be preferred to a greater future good’ (allowing for difference of certainty).” Rawls (1971:293) reiterates the point: “A present or near future advantage may be counted more heavily on account of its greater certainty or probability, and we should take into consideration how our situation and capacity for particular enjoyments will change. But none of these things justifies our preferring a lesser present to a greater future good simply because of its nearer temporal position.” While discount rates feature prominently in economic models of decision making, there is an important difference between the discounting of positive and negative experiences—which John Broome (1991) calls “pure” discounting—and the discounting of commodities used by economists in cost-benefit analysis. This paper concerns pure discounting.

status of past near bias, insofar as philosophers have often attempted to argue for normative conclusions from the presence of certain descriptive matters of fact.²

There are various hypotheses one could have regarding our preferences over past events. First, one could think that there will be no past near bias. That is, one could hypothesise that all near bias is future-directed, and that people do not care about the nearness of past events at all. Second, one could hypothesise that people will have preferences regarding the nearness of past events, but that these preferences will be entirely different from their preferences for future events. Third, one might predict that people will show symmetrical preferences: that their preferences for past events will mirror their preferences for future events, and hence that insofar as individuals are future near biased, they will be past near biased.

While little consideration has been given either to the descriptive or normative aspects of past-directed near-biased preferences, it is easy to see that one could appeal to descriptive facts about past-directed preferences in the service of arguments about the normative status of future-directed preferences. For instance, if the first hypothesis were true—if people do not display past near bias—then that might support the view that future near bias is irrational. One could argue that in such a case we are treating the temporal directions differently, and that this treatment is unmotivated or ad hoc. If future-directed and past-directed preferences ought to be similar (in this regard) then one set of preferences is irrational. Extant research into future near bias could then be used as reason to think that it is future-directed, and not past-directed, preferences that are most likely to be irrational. At the very least, if the first hypothesis is borne out, proponents of the rationality of future near bias either need to argue that our past-directed preferences are irrational, or they need to explain why the temporal asymmetry of our preferences is rationally permissible.

By contrast, if the third hypothesis is true—if people show symmetric past and future near bias—then this would provide some support for those who defend the permissibility of future near bias. Indeed, what little empirical evidence there is about past-directed preferences seems to suggest that patterns of past near bias will mirror patterns of future near bias (Yi et al. 2006; Bickel et al. 2008). This descriptive finding would not, of course, show that this pattern of preferences is rational, but it would undermine any argument that attempted to show that future near bias is irrational by pointing to some sort of deep difference between our future- and past-directed preferences. Finally, exactly what implications a vindication of the second hypothesis would have for the normative status of future near bias is unclear. That, presumably, would depend on exactly what sorts of past-directed preferences are found.

The upshot is that we cannot even begin to consider these arguments without first investigating past-directed preferences.

This paper aims to experimentally test the third hypothesis just mentioned: that insofar as we find certain patterns of future near bias, so too we will find the same patterns of past near bias. We call this the *past/future symmetry hypothesis*. We suppose that insofar as there is such a symmetry between past and future near bias, this is the result of individuals showing the same bias (or lack thereof) in the past, as they do in the future. That is, insofar as an individual is future near-biased, we expect that they will also be past near-biased, and insofar as they are future time-neutral we expect they will also be past time-neutral.

Further, following recent research (Greene et al, ms.) we expected to find a difference in participants' preferences regarding the temporal location of hedonic events versus non-hedonic events. *Non-hedonic* events are those that do not concern pleasurable

² Examples of philosophers who argue against the rationality of future near bias include Sidgwick (1884), Lewis (1946), Elster (1970), Nagel (1970), Rawls (1971), Broome (1991), and Brink (2011).

or painful experiences. These are typically events that are not experienced directly by the agent, or where the agent's experience is not tied to the temporal location of the event. In particular, Greene et al. found participants to be *future biased* for *hedonic events*. In other words, participants prefer pleasurable experiences to be in their future rather than their past, and painful experiences to be in their past rather than their future. However, for *non-hedonic* events, they found participants to be time neutral. In other words, participants had no preference regarding whether positive or negative non-hedonic events were in their past or future. Hence, we hypothesise that the same pattern will be seen with regards to non-hedonic near bias: people will be indifferent to where, in the near or far future and the near or far past, *non-hedonic* positive and negative events are located.

In sum, then, we hypothesise that we will find the same pattern of past-directed preferences that we see for future-directed preferences. When the events in question are hedonic, we expect participants to be both future and past near-biased, and when the events are non-hedonic, we expect to see both future and past time-neutrality. In §2, we summarise the existing experimental research on hedonic future near bias. Then, in §3 we present our experiments and results. In §4 we discuss those results and their relevance to philosophical inquiry on near bias. There, we note that our results support our past/future symmetry hypothesis. In this, we replicate the findings of Yi et al. (2006) and Bickel et al. (2008). We also add to Greene et al's (ms) finding that time-biased preferences are different for hedonic and non-hedonic events. We found that participants were near biased for hedonic events but had no such preference for non-hedonic events.

There are, however, two surprising outcomes to this research. First, we found a good deal of variation in the population regarding these preferences. That is, for almost every preference that can be had with regard to future events, there is some substantial sub-population that expressed that preference; *mutatis mutandis* for preferences over past events. Second, although there is past/future symmetry of preferences at the population level, this does not appear to be the result of individual past/future symmetries. That is, we found no correlation between an individual's preferences over future events and their preferences over past events. This, which appears to contradict findings of Yi et al. (2006) and Bikel et al. (2008), suggests that the mechanisms responsible for our preferences over past events are different from those responsible for our preferences over future events.

Our population-level results seem to undermine certain forms of argument for the irrationality of future near bias—especially those that appeal to past/future asymmetry of preferences—since these results show there to be past/future symmetry of preferences. In fact, though, matters are more complex than this, given that our results suggest that the mechanisms that generate future preferences are different from those that generate past preferences. While that observation does not show that either set of preferences is irrational, it could lend support to arguments against near bias that are based on the fact that individuals appear to treat the two temporal directions differently.

2. Prior Research on Near Bias

Our hypothesis is that we will see the same pattern of past-directed preferences as we do future-directed preferences; i.e., we will find past/future symmetry. Our reason to suppose we will find this symmetry is based on the only two studies of past-directed near bias that have been undertaken. Both of these studies, the first by Yi et al. (2006) and the second by Bickel et al. (2008) required participants to make various choices about when, and how much, (hypothetical) money they would receive in the future or did receive in the past. In both studies a discount rate was calculated for each participant regarding both their future- and past-directed preferences. These studies have two main findings

that are of interest in this context. First, they both found that future- and past-directed preferences were symmetrical: insofar as participants were inclined to future discount, they were also inclined to past discount. Second, both studies found correlations between individual participants' future discount rate and their past discount rate across a range of conditions.

Since philosophers are interested in the rationality of preferences over all kinds of events—not just monetary gains or losses—in this paper we focus on preferences over both hedonic and non-hedonic events. While there is no research we know of that looks at near bias over both of these kinds of contexts, there is research that focuses more broadly on past- and future-directed attitudes to hedonic and non-hedonic events (see Caruso et al. (2008); Caruso (2010); Burns et al. (2012) and Greene et al. (ms)). Given this research, what should we expect the hypothesised symmetry between past-directed and future-directed preferences to look like?

Philosophers typically assume that people are hedonically future near-biased.³ They base this assumption on work in the social sciences which, taken as a whole, seems to support this view. Early experiments on temporal discounting indicate a “pervasive devaluation of the future” (Ainslie and Haslam (1992:59)). For example, in the experiments of Thaler (1981), Hausman (1979), and Akerlof, (1991), people assigned less value to future money, time, and effort, respectively, than their present analogues.⁴ Similar results in animal studies backed this idea; e.g., Green et al. (1981) demonstrated temporal discounting in pigeons choosing between food rewards. These results suggested a widespread and robust phenomenon.⁵ However, in a meta-analysis of studies done between 1978 and 2002, Frederick et al. (2002:377) found “tremendous variability” in estimates of people’s average discount rate. They show that existing experiments estimate an annual discount rate between -6%—where negative discounting rates are, in effect, rates according to which goods devalue as they become closer to one’s location—and infinity, with no clustering of moderate results.

Frederick et al. (2002) and Soman et al. (2005) conclude that researchers should understand intertemporal choices as the *joint product* of *conflicting* motives.⁶ Most importantly, many experiments demonstrate that people sometimes prefer to incur a loss immediately rather than delay it.⁷ One explanation for this is that delaying an unpleasant experience results in disutility from anticipation (i.e., negative “waiting”). Other experiments suggest that people often prefer improving sequences over declining ones

³ For more detailed definitions of hedonic future near bias see Greene and Sullivan (2015:949)

⁴ For reviews, see Ainslie and Haslam (1992); Frederick et al (2002); and Soman et al. (2005).

⁵ Historically, the dominant model in social science is the *discounted-utility model*, according to which all of the underlying motives involved in intertemporal choice can be condensed into a single parameter: the discount rate (Frederick et al., (2002:351)). The discount rate defines precisely how much agents discount the value of a good based on its temporal distance. To determine discount rates, experimenters typically proceed by presenting subjects with a choice between sooner smaller and larger later rewards, and they then combine the responses into an average discount rate. A second wave of research focused on *dynamic inconsistencies* in temporal discounting. Specifically, experiments suggested that subjects often prefer a smaller sooner reward to a larger later one when the smaller reward is available immediately, but that this preference flips when an equal delay is added to both rewards. When interpreted within the framework of the discounted-utility model, this suggests that people engage in *hyperbolic*, rather than *exponential* discounting. For hyperbolic discounters, valuations fall rapidly for near delays and slowly for distant delays. See Ainslie and Haslam (1992); Hoch and Loewenstein (1991); Read and van Leeuwen (1998).

⁶ As Loewenstein (1987:666–7) and Frederick et al. (2002:352–3) show, this was the predominant view of economists before the introduction of the discounted-utility model in 1937. Frederick et al. write, “When the [discounted-utility] model eventually became entrenched as the dominant theoretical framework for modelling intertemporal choice, it was due largely to its simplicity and its resemblance to the familiar compound interest formula, and not as a result of empirical research demonstrating its validity.”

⁷ See Benzion, Rapoport, and Yagil (1989); Loewenstein (1987); and MacKeigan et al. (1993).

(Loewenstein and Prelec, 1991), or to maximize the utility of anticipation by delaying positive events (i.e., positive “waiting”) (Loewenstein 1987; 1996).

In sum, while there is strong evidence suggesting that people devalue distant-future events in comparison to their near-future analogues, this devaluation is countered by considerations from the utility of anticipation (positive waiting), the disutility of anticipation (negative waiting), and preferences for improving sequences.

Given this, while we expected to find hedonic future near bias amongst our participants—and hence, given our past/future symmetry hypothesis, we expected to find hedonic past near bias—we also expected to find sub-populations that would show the reverse: those who preferred to delay hedonic events to later times due to the utility of positive waiting, and those who preferred not to delay negative hedonic events due to the disutility of negative waiting. Since we predicted past/future symmetry, then, we also predicted that we would find analogous patterns with regard to past-directed preferences.

There has been little research into non-hedonic future near bias, and none on non-hedonic past near bias. Given the results found by Greene et al. (ms), we expected to find non-hedonic time neutrality both towards the future and the past. Again, however, we expected to find substantial sub-populations with different preferences regarding non-hedonic events (both in the past and in the future).

In what follows we test these predictions.

3. Experimental Design and Results

3.1 Method

3.1.1 Participants

534 people participated in the study. Participants were U.S. residents, recruited and tested online using Amazon Mechanical Turk and Qualtrics, and compensated \$0.50 for approximately 5 minutes of their time. 46 participants had to be excluded for failing to follow task instructions. This means that they failed to answer the questions (40), or failed the attentional check question (6). The remaining sample was composed of 488 participants (aged 17-83; 193 female; 4 preferred not to answer). Mean age 36.8 ($SD = 12$). Participants were recruited and tested the same way in all experiments reported. Ethics approval for this study was obtained from the [blanked] Human Research Ethics Committee. Informed consent was obtained from all participants prior to testing.

3.1.2 Materials and Procedure

Participants were randomly assigned to one of four conditions: positive or negative valence; hedonic or non-hedonic event, in a mixed within/between-subjects design. These four conditions reflected all possible combinations of valence (positive or negative) and kind of event (hedonic or non-hedonic). Hence we tested the following four conditions: hedonic positive valence; hedonic negative valence; non-hedonic positive valence; non-hedonic negative valence. Each participant was assigned to only one condition.

We developed a single base vignette that could be minimally modified for each of these conditions. Participants either read a vignette describing a positively valenced non-hedonic event (receiving a community award) or a negatively valenced non-hedonic event (having embarrassing photographs released), or a positively valenced hedonic event (receiving a most favoured meal) or a negatively valenced hedonic event (receiving a most disfavoured meal).

The vignettes are below. The first is the non-hedonic base vignette with both positively and negatively valenced substitutions listed. The second is the hedonic base vignette with both positively and negatively valenced substitutions listed.

Imagine you are an astronaut on a 10-year voyage from Earth to set up a colony on a new planet. It is a one-way mission, and there is no way you can return to Earth. You are 5 years into the voyage. Just before you left, you learned that [your home-town mayor plans to award you an important community service prize]/[someone plans to release embarrassing photos of you] at some time during the 10-year period in which you are traveling. You do not know when they will [award the prize]/[release the photos], and it is not possible to communicate with Earth during the trip, or even once you have arrived on the new planet. You find yourself wondering whether [the prize has been awarded]/[photos have been released] yet.

Imagine you are an astronaut on a 10-year voyage between planets. 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 meal]/[most disliked meal] — [which you really like]/[which you really 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.

We developed these vignettes to avoid potential confounds arising due to participants having preferences over the location of events that do not reflect any pure time preference. For instance, they might reflect the fact that participants reason that in time, others will forget negatively valenced events—the embarrassing photos—(which is desirable) and positively valenced events—the community award—(which is not desirable).⁸ To avoid this confound the vignette describes a one-way trip from Earth to another planet, with no communication between Earth and the ship.

In each condition, participants were asked how much they agree, on a Likert scale from 1 (strongly disagree) to 7 (strongly agree) with two randomly selected statements.

Participants in the non-hedonic conditions were asked how much they would agree with *one* of the following statements:

(a) I would prefer to learn that the [important community service prize]/[embarrassing photos] will be [awarded]/[released] tomorrow, and will not be [awarded]/[released] 1 year into the future.

(b) I would prefer to learn that the [important community service prize]/[embarrassing photos] will be [awarded]/[released] 1 year into the future, and will not be [awarded]/[released] tomorrow.

And *one* of the following statements:

(c) I would prefer to learn that the [important community service prize]/[embarrassing photos] [was awarded]/[were released] yesterday, and [was not awarded]/[were not released] 1 year in the past.

⁸ Cf. Greene (ms),

(d) I would prefer to learn that the [important community service prize]/[embarrassing photos] [was awarded]/[were released] 1 year in the past, and [was not awarded]/[were not released] yesterday.

Participants in the hedonic conditions were asked how much they would agree with *one* of the following statements:

(a) I would prefer to learn that my [favourite]/[most disliked] meal will be dispensed tomorrow, and will not be dispensed 1 year into the future.

(b) I would prefer to learn that my [favourite]/[most disliked] meal will be dispensed 1 year into the future, and will not be dispensed tomorrow.

And *one* of the following statements:

(c) I would prefer to learn that my [favourite]/[most disliked] meal was dispensed yesterday, and was not dispensed 1 year in the past.

(d) I would prefer to learn that my [favourite]/[most disliked] meal was dispensed 1 year in the past, and was not dispensed yesterday.

After responding to each statement, participants were then asked to indicate their level of confidence in their judgement. After having done so for both statements participants proceeded to a new screen, where they answered a comprehension question: “*In this vignette, you were asked to imagine that you were...*” with the answer options (1) an astronaut (2) a dog. Participants who chose (2) were excluded. At no point could participants return to a previous screen.

3.2 Results

Before reporting the statistics and details, first we here provide a summary of the findings. We found that participants preferred hedonic events to be nearer to them in time than non-hedonic events, and that they preferred positively valenced events to be nearer to them than negatively valenced events.

We found similar average judgments for the past as for the future. However, there was a high degree of heterogeneity, with a large proportion of people that did not show future near bias or past near bias.

At the individual participant level, we had hypothesised that participants who showed a particular bias for the future, or not, would show the same bias (or lack of bias) for the past. In other words, we expected participants’ judgments would show temporal symmetry. Instead, we found no correlation between participants’ judgments about the past and their judgments about the future.

In order to control for effects of the wording of the questions, in each condition participants were asked two of four question types. For example, half of the participants in the positive hedonic condition were asked how much they *agreed* with the statement that they would prefer their favourite meal to be dispensed tomorrow, and not one year in the future, while the other half were asked how much they *agreed* with the statement that they would prefer their favourite meal to be dispensed one year in the future, and not tomorrow. To appropriately combine the results from these two forms of the questions, levels of agreement with the latter question were reverse-coded (i.e. a response of 1 was transformed into a response of 7; a response of 2 was transformed into a response of 6, and so on). After this reverse coding, the results are *as if* all participants had been asked their levels of agreement that they would prefer the event in question be

located tomorrow, and not 1 year in the future. More generally, in what follows, *higher levels of agreement* indicate participants' collective preference for the event in question to be located in the *near* past or future (whichever is relevant).

Level of agreement and level of confidence judgments were analysed using separate repeated-measures ANOVAs. The ANOVAs included a within-subjects factor of temporal direction (past; future) and between-subjects factors of valence (negative; positive) and event type (hedonic; non-hedonic). We also calculated Pearson correlation coefficients to assess the association between levels of agreement for the past with that of the future.

Level of Agreement

The 2x2x2 repeated-measures ANOVA revealed a main effect of event type $F(1, 484) = 7.122, p = .008$ and valence $F(1, 484) = 23.003, p < .001$. The third factor, direction, did not have a significant effect, $F(1,484) = 1.660, p = .198$. No significant interaction effects were observed.

The main effect of event type showed that levels of agreement were significantly higher for hedonic events ($M = 4.17, SD = 1.45$) than for non-hedonic events ($M = 3.82, SD = 1.46$). That is, participants overall preferred hedonic events to be nearer to them in time than non-hedonic events.

The main effect of valence showed that levels of agreement were significantly higher for positively valenced events ($M = 4.31, SD = 1.46$) than for negatively valenced events ($M = 3.68, SD = 1.45$). That is, participants overall preferred positively valenced events to be nearer to them than negatively valenced events.

There was no main effect of direction (past versus future), suggesting that participants overall behaved similarly in the past and the future. Critically, however, little to no correlation was evident between individual participants' responses for the questions about the past and those about the future. The Pearson correlation coefficient calculated between levels of agreement in the past conditions and levels of agreement in the future conditions was not statistically significant ($r = .023, t(486), p = .607$) and the narrow 95% confidence interval (-0.067, 0.112) suggests there is in fact close to no relationship. Breaking the data down by the four conditions (negative hedonic, negative non-hedonic, positive hedonic, positive non-hedonic), the estimated correlations are similarly close to zero ($r_s = -.008, .093, -.017, -.118$) and not statistically significant ($p_s = 0.933, 0.315, 0.850, 0.196$). Hence, the overall apparent symmetry between past and future was not representative of individuals' preferences.

When presented with certain sorts of questions, participants give a restricted range of Likert scale responses, and restriction of range can result in very small correlations. Here, however, while indifference was sometimes the modal response, participants did use the entire range of responses (see Figure 1). Substantial numbers of people strongly preferred, strongly did not prefer, and indicated they were completely neutral on the questions in this study. This may well reflect distinct groups of people in the population, but getting strong support for that would require further investigation.

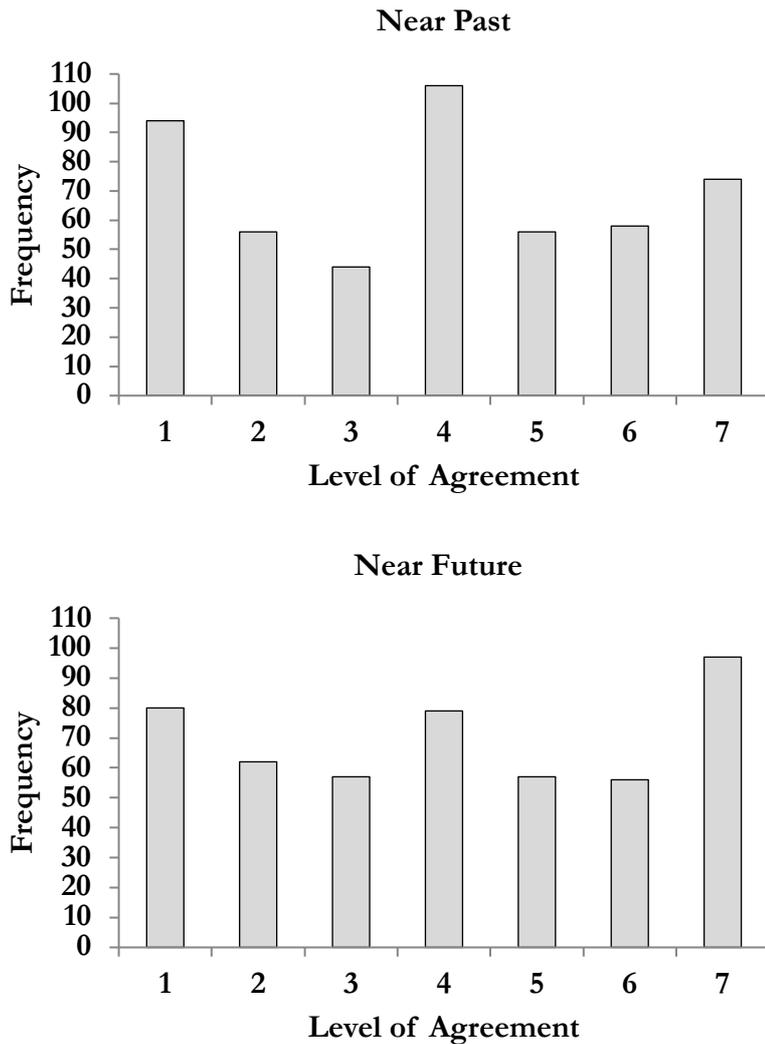


Figure 1. Level of agreement Judgments for past (top) and future (bottom) conditions, collapsed across hedonic and non-hedonic events and positive and negative valence. “1” is strongly disagree, “7” strongly agree with a statement in the direction that indicates a near preference.

The pattern of data underlying the net zero or close to zero correlation within participants is shown in Figure 2. The lack of a statistically significant correlation is consistent with the proposition that a participant’s preference for near or remote events in the past provides no information about whether they prefer an event to be near or remote in the future.

Although there is no correlation between participants’ past- and future-directed preferences, participants do appear to differ in the *magnitude* of the responses they tend to give: that is, the same participants who have more extreme preferences about the past, tend to have more extreme preferences about the future. This could account for the apparent clustering of responses at the corners and in the centre—consider that if all participants gave the same magnitude of answers for the past questions as for the future questions, the data would form an “X” in Figure 2. To explore this possibility, we calculated the magnitude of participants’ answers, subtracting 4 from every response and then taking the absolute value.

The magnitude scores were positively correlated for past and future, suggesting that participants’ responses to different questions do tend to be similar in magnitude. For negative hedonic events $r = 0.252$ ($t(123) = 2.893$, $p = 0.005$), negative non-hedonic $r =$

0.546 ($t(116) = 7.014$, $df = 116$, $p < .001$), positive hedonic $r = 0.394$ ($t(121) = 4.709$, $p < .001$), and positive non-hedonic $r = 0.626$ ($t(120) = 8.784$, $p < 0.001$).

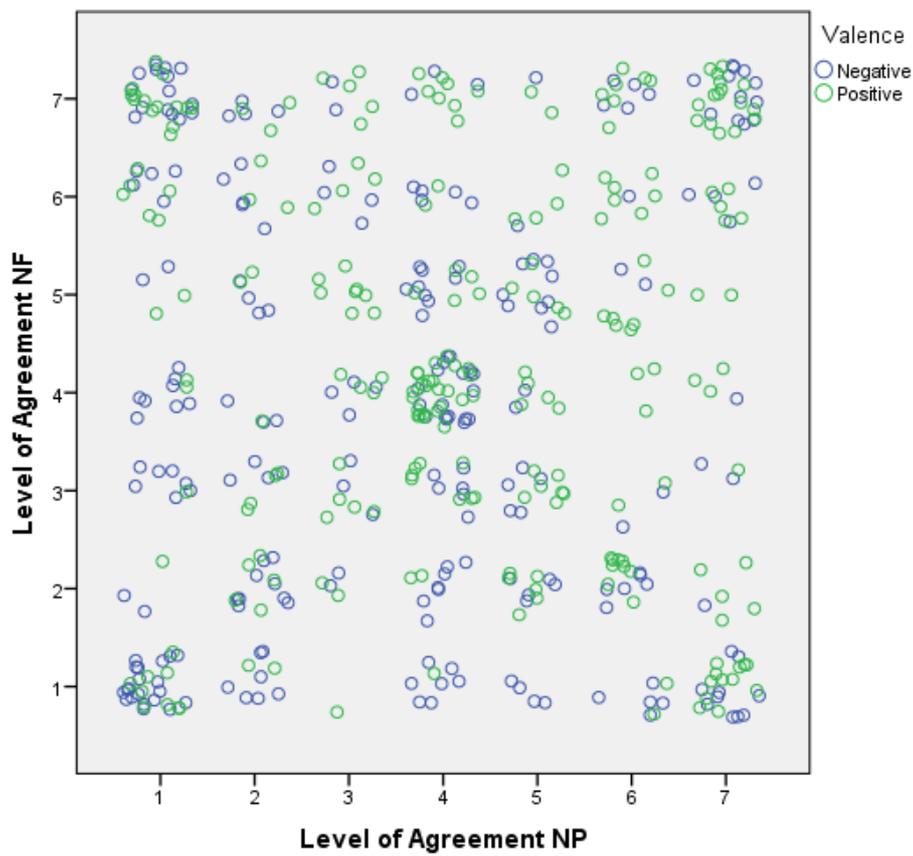
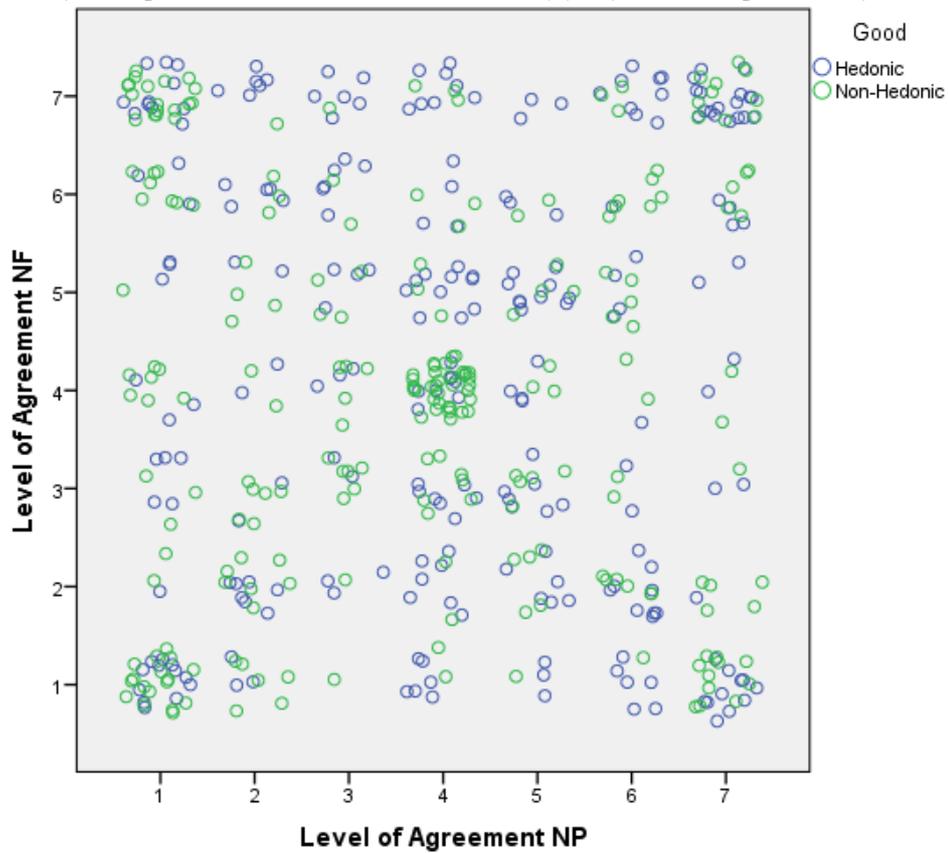


Figure 2. Individual participant responses. Preference for near past against preference for near future. Top, data is broken down by hedonic (blue) and non-hedonic (green). Bottom, data is broken down by negative valence (blue) and positive valence (green). No association is evident between the near-past and near-future levels of agreement (see text). Points are jittered to avoid complete overlap.

Level of Confidence

The 2x2x2 repeated-measures ANOVA found a main effect of direction $F(1, 484) = 5.447, p = .02$ and a two-way interaction between direction and event $F(1, 484) = 4.968, p = .026$. No other significant main effects or interaction effects were observed.

The main effect of direction was that confidence was slightly higher for the future conditions ($M = 5.59, SD = 1.46$) than for past conditions ($M = 5.437, SD = 1.59$).

Simple effects tests using a Bonferroni correction for four comparisons were carried out on the two-way interaction between direction and event. First, for past judgments there was no significant difference in levels of confidence between hedonic events ($M = 5.48, SD = 1.57$) and non-hedonic events ($M = 5.40, SD = 1.58; p = .558$). Second, for future judgments, levels of confidence were significantly higher for hedonic events ($M = 5.78, SD = 1.46$) than non-hedonic events ($M = 5.40, SD = 1.46; p = .005$). Third, for hedonic events, levels of confidence were significantly higher for future judgements than for past judgments ($p = .001$). Fourth, for non-hedonic events there was no significant difference in levels of confidence between past and future judgments ($p = .941$; see Figure 3).

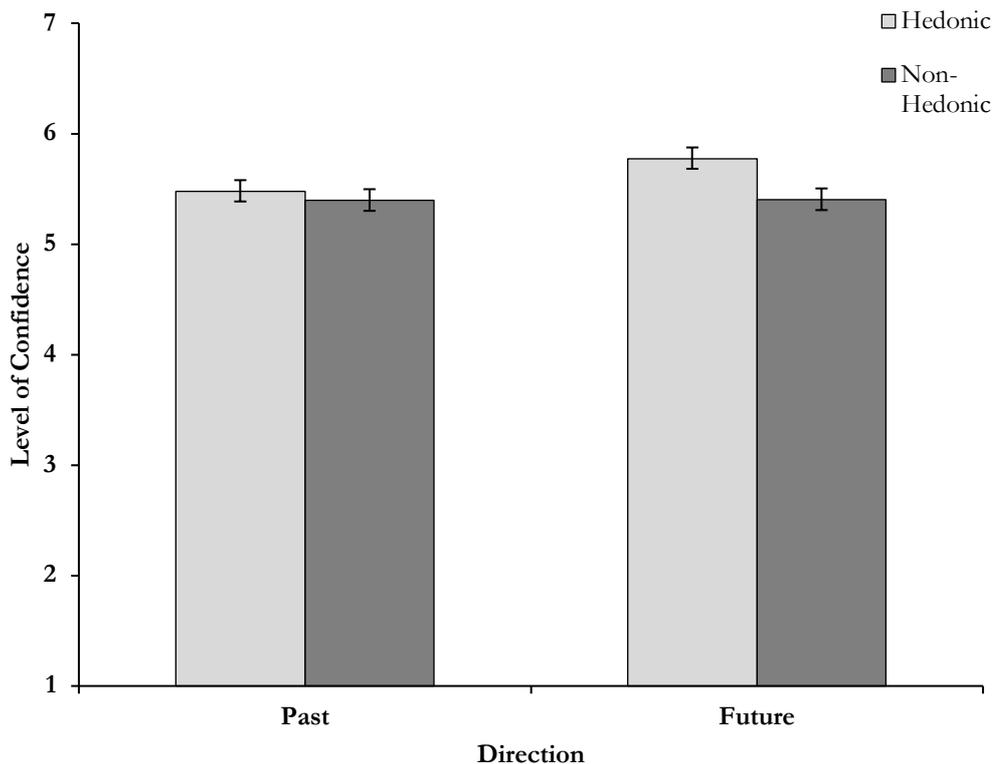


Figure 3. Levels of confidence, collapsed across valence. The main effect of direction was significant, as was the interaction between direction and event Error bars are one standard error.

Participants' level of confidence in their response to their past condition was positively correlated with that of their future condition judgment, $r = .558, t(486) = 14.81, p < .001$.

That is, the more confident someone was in their judgment in one condition, the more confident they were in the other condition. This is much like the magnitude of participants' preferences, which were also correlated (see previous section). In fact it is possible that the confidence correlation is a consequence of the correlation in preference magnitudes. Suchman (1950), among others, has observed that stronger preferences are correlated with higher certainty.

4. Discussion

Before we discuss these results, it's worth noting a potential weakness of this study. We use hypothetical scenarios to test participants' preferences. By contrast, at least in some behavioural economics studies, participants' preferences are tested under non-hypothetical conditions in which, for instance, they gain or miss out on real money. Clearly, the testing of past-directed preferences cannot be done in anything other than a hypothetical manner. Nevertheless, it's worth noting that a number of studies have found very similar results when comparing participants' responses to hypothetical and non-hypothetical scenarios (Lagorio and Madden, 2005; Johnson and Bickel (2002)).

Our data shows population-level symmetries between past- and future-directed preferences regardless of whether the events are hedonic or non-hedonic, and regardless of their valence. In this regard, our findings were as we predicted based on the earlier findings of Yi et al. (2006) and Bickel et al. (2008). Also, as found in previous studies, on average participants showed both past and future near bias.

An important caveat to these observed differences in average ratings among conditions is the high heterogeneity of the participants' judgments. The preferences expressed were extremely variable, with some indication of trimodality—that is, peaks of roughly similar height at both extremes and at the indifference point. This suggests that it may be misleading to interpret the data via the average ratings. Likert ratings nearly always contain some variability, which ordinarily may be attributed to noisy responses where the source of noise is orthogonal to the issues of interest, or is due to a population comprising people that differ in how they use rating scales, and may have different preferences. In such situations there can be a meaningful central tendency or majority group. Here, however, the variability is large enough that one should consider the possibility that there are multiple sub-populations of approximately equal size having very different preferences. Separating out the sub-populations is likely to be difficult even with large numbers and sophisticated statistics.

Although on average, the participants showed near bias with regard to positive hedonic events (both in the past and future), a large proportion of people did not express any preference or expressed the opposite preference. Moreover, we see fewer near biased preferences, on average, with regard to non-hedonic events (positive and negative) than hedonic events (positive and negative).

In further contradiction to what was expected before data collection, there was little to no correlation between individuals' future-directed and past-directed preferences. This is especially puzzling, particularly given that Yi et al. (2006) and Bickel et al. (2008) both found correlations between past-directed and future-directed preferences. In light of this, two things call out for explanation. The first is why there is such a difference between these results and those of the Yi et al. and Bickel et al. studies. The second is why, given that we found no such correlations, we nevertheless find population-level symmetries. We will consider each of these in turn.

First, there are some notable differences between the Yi et al. and Bickel et al. studies and this study. As noted previously, both of those studies investigated preferences over (hypothetical) monetary gains and losses in the future and past. One might expect there to be more uniformity in past-directed and future-directed

preferences over monetary gains and losses than over various kinds of hedonic and non-hedonic events. Moreover, both these studies used the preferences of participants to determine a discount rate for both the past and future. In our study, however, we aggregated strength of preferences on a Likert scale. Given that the studies all appeal to mean or median values, however, and given the very high degrees of heterogeneity we found amongst our (much larger) sample size, we are not well placed to evaluate why these studies found correlations that we did not. We think a primary focus of future work should involve finding less variable sub-populations and investigating their preferences. Nevertheless, in what follows we offer some hypotheses about what our results suggest.

Our suggestions take their cue from Frederick et al. (2002) and Soman et al.'s (2005) hypothesis that intertemporal choices are the joint product of conflicting motives (discussed in Section 2). From this perspective, to understand intertemporal choice we must adopt a psychological approach that seeks to understand the numerous and disparate motives underlying intertemporal choice, and resist the impulse to condense all these factors into a discount rate. Just as factors other than the temporal distance and valence of an event can affect people's future-directed preferences—factors such as negative waiting and positive anticipation—so too do additional, but, crucially, *different* factors, influence people's past-directed preferences.

To further demonstrate the plausibility of a psychological approach to future and past near-biased preferences, in what follows we offer some possible explanations of the factors behind respondents' choices in each condition.

Consider hedonic events first. We predicted that participants would prefer to have their least favourite meal a year ago, and to have their most favourite meal a day ago (rather than the other way around). Why might some participants have responded differently? Some participants might prefer to have their least favourite meal yesterday because they find *comparisons* salient. For a wide range of mental processes, beginning with sensory systems and extending to cognition, humans and animals tend to use relative (comparative) values rather than the absolute values of things (e.g. Tversky & Kahneman, 1974; Festinger 1954; Khaw, Glimcher, & Louie, 2017; Wark, Lundstrom, & Fairhall, 2007). In particular, there is evidence that upward comparisons (those comparing oneself to someone better off in some relevant respect) can lead to increased negative emotions, including jealousy (Salovey & Rodin (1984)) and frustration (Martin, 1986), and to lowered self-evaluations (Marsh & Parker (1984)), which decrease happiness (Pleban & Tesser (1981)). By contrast, downward comparisons (those comparing oneself to someone worse off in some relevant respect) typically enhance well-being (Wills (1981); Wood, Taylor & Lichtman (1985)). This suggests that there are reasons for individuals to prefer that events are located and arranged in such a way as to support certain kinds of comparative judgements, and to avoid others.

For instance, if my least favourite meal was yesterday, then today's bland meal will, by comparison, seem quite tasty. By contrast, if my most favoured meal was yesterday rather than a year ago, then today's bland meal will, by comparison, seem extra bland. Hence there are reasons arising from comparisons that might lead individuals to have the opposite preferences to those displayed by past near-biased individuals. Individuals motivated by comparisons want to maximise the tastiness of *today's* meal, and so they have preferences over the temporal locations of the least favourite, and most favourite, meals, which are designed to do so. Thus, individuals have a reason to prefer their least favourite meal being located yesterday and their most favourite meal being located a year ago that have nothing to do with past-directed analogues of positive and negative waiting. Thus, if certain kinds of comparisons are salient to participants, we

would expect them to prefer to have their most favourite meal located a year ago, and their most disliked meal a day ago.

Why might someone, instead, prefer to have the most favourite meal yesterday, and the least favourite meal one year ago? That is, why might someone have past near-biased preferences? Here, we think, considerations arising from memory might explain these preferences. Research shows that most memories fade over time: memories for recent events typically contain more sensorial and contextual detail than those for remote events (Johnson et al., (1988)), while memories of positive experiences often contain more such details than memories of negative experiences (Byrne, Hyman, & Scott (2001); D' Argenbeau, Comblain, & Van der Linden (2003); Destun & Kuiper (1999)). Hence, if one wants to maximise the extent to which one remembers a pleasant event, one may prefer that event to be in the near past, rather than the far past. Likewise, if one wants to minimise the extent to which one remembers an unpleasant event, one may prefer that event to be in the far past, rather than the near past. On the assumption that people generally want to remember pleasant events and forget unpleasant events, it makes sense for people to be past near biased. Thus, if these memory effects are salient to participants, we would expect them to prefer to have their most favourite meal located yesterday, and their most disliked meal a year ago.

We will call the first of these factors *the comparison effect*, and the second *the memory effect*. The comparison effect and the memory effect pull in different directions when it comes to hedonic events. The memory effect will push individuals to prefer to have positively valenced hedonic events in the near past and negatively valenced hedonic events in the far past. The comparison effect will push individuals to prefer to have positively valenced hedonic events in the far past and negatively valenced hedonic events in the near past. Given this, it might be that those for whom comparisons are especially salient will prefer to have negative hedonic events in their near past and positive hedonic events a year ago. Those for whom the memory effect is especially salient will prefer to have negative hedonic events located a year ago and positive hedonic events yesterday. Those for whom neither of these are especially salient, or for whom they are equally salient, may well be indifferent between the temporal locations of these events, assuming there are no other factors at play.

This is exactly what we find. When asked their preference regarding where hedonic positively valenced events are located, 17% were indifferent, 37% strongly preferred, or preferred, that the event be located in the near past, and 20% strongly preferred, or preferred, that the event be located in the far past. When asked their preference regarding where hedonic negatively valenced events are located, 24% were indifferent, 21% strongly preferred, or preferred, that the event be located in the near past, and 32% strongly preferred, or preferred, that the event be located in the far past.

Moreover, if something like this is the right explanation for this pattern of preferences, it is easy to see why we find no correlation between whether an individual is future near biased and whether they are past near biased. For the mechanisms underlying the two kinds of bias are entirely different, and, plausibly, independent of one another.

Turn now to participants' preferences about non-hedonic events. As we predicted, we find time-neutral preferences for both positive and negative non-hedonic events. However, we also found substantial sub-populations that had time-biased preferences with regard to both future and past events. One might think it especially puzzling that anyone would have past-directed preferences regarding the temporal location of non-hedonic events. Yet we find 46.6% of our participants strongly prefer, or prefer, to have negatively valenced non-hedonic events located a year ago while 21.2% strongly prefer, or prefer, to have negatively valenced non-hedonic events located

yesterday. Similarly, 24.6% strongly prefer, or prefer, to have positively valenced non-hedonic events located a year ago and 28.7% to have them located yesterday

Once again, this data suggests that there are a number of factors that are being taken into consideration in the determination of past-directed preferences. It has long been hypothesised that individuals care not just about the total utility of their lives, but about the *distribution* of that utility. Indeed, some philosophers, such as Velleman (1991) and Slote (1983:23-24) hold not only that people have such preferences, that but they are rational to do so. Both suppose that an improving life is better than a worsening life even if the 'area under the curve' is the same; i.e., even if the aggregate utility is the same. Our suggestion is that some combination of preferences regarding the shape of one's life, and preferences regarding the aggregate utility in one's life, jointly serve to explain the pattern of responses that we found regarding the location of non-hedonic events, and particularly non-hedonic events in the past.

First, if some participants have a preference for an improving life, then this would explain a preference for negatively valenced events to be maximally past and positively valenced events to be maximally future. When it comes to future-directed preferences this manifests as a preference for negative events to be in the near future and positive events in the far future (a preference that can also be explained by positive and negative waiting). When it comes to past-directed preferences this will manifest as past near bias. A preference to have an improving life will likely be reflected more in preferences over non-hedonic events, such as the release of embarrassing photos or the receipt of a community service award, than it will over fleeting events such as the quality of a meal.

One way in which 'shape of life' considerations might be particularly salient in the context of this study lies in the 'shape' or distribution of reputational factors. The non-hedonic events in this study both involve reputational factors: the award of an important community service prize and the release of embarrassing photos. Recent research suggests that reputation preservation is important within social contexts because building a positive reputation enhances evolutionary fitness (Alexander (1987); Nowak & Sigmund (1998); Nowak & Sigmund (2005)) and explains prosocial behaviour (Rabin (1993); Dufwenberg & Kirschsteiger (2004)) by signalling trustworthiness.

Plausibly, however, participants might have different preferences about the shape of reputational factors. Some might be especially sensitive to *current* reputational effects (i.e., effects to their reputation at the time they are asked about their preferences), while others might be especially sensitive to *improvements* in reputational factors over time. In either case, such participants will prefer to have the positive non-hedonic event located yesterday and the negative non-hedonic event located a year ago. By contrast, other participants might be more sensitive to aggregate reputational factors rather than to those factors' distribution. These participants will prefer to have the negative non-hedonic event located yesterday and the positive non-hedonic event located a year ago, since these preferences result in there being *more* years of positive reputation and *fewer* years of negative reputation, and hence more aggregate positive over negative reputation.

Of course, other factors might also be at play. Our key claim is that there is a range of factors that determine past-directed preferences, and that at least some of these are different from those which determine future-directed preferences. This explains why there is no correlation between individual future-directed preferences and past-directed preferences. Second, it suggests that different factors are most salient for different individuals, thus explaining why different individuals show rather different future- and past-directed preferences. Third, it suggests a number of interesting connections between normative arguments and descriptive findings.

Our findings undermine any argument for the rationality of future near bias based on symmetries with past-directed preferences. Although there are population-level symmetries, these are not the result of individual-level symmetries. Our findings also undermine arguments which attempt to show that future near bias is irrational based on the *absence* of past near bias. For we found that people are as likely to be past near-biased as they are future near-biased. Finally, our findings tend to undermine any *general* claims about the normative status of future near bias that appeal to the idea that there is a single, shared, future-directed set of preferences (or a single, shared, past-directed set of preferences). For in fact, we find substantial sub-populations who are ‘far-biased’ (both in the past and the future) and substantial sub-populations that are temporally neutral (both in the past and the future). That descriptive fact tells us nothing about which of these attitudes is rationally permissible or required. But it does tell us that we cannot infer that one is permissible, or required, based on some actual descriptive convergence in preferences—for there is much less convergence than one might have supposed.

References

- Ainslie, G., & Haslam, N. (1992). Self-control. *Choice over time*, 177, 209.
- Akerlof, G. A. (1991). Procrastination and obedience. *The American Economic Review*, 81(2), 1-19.
- Alexander R. (1987). *The biology of moral systems*. New York, NY: Aldine de Gruyter.
- Benzion, U., Rapoport, A., & Yagil, J. (1989). Discount rates inferred from decisions: An experimental study. *Management science*, 35(3), 270-284.
- Bickel, W. K., Yi, R., Kowell, B. P., and Gatchalian, K. M (2008). “Cigarette smokers discount past and future rewards symmetrically and more than controls: Is discounting a measure of impulsivity”. *Drug and Alcohol Dependence* 96: 256–262.
- Brink, D. O. (2011). “Prospects for Temporal Neutrality”. in Craig Callender (ed). *The Oxford Handbook of Philosophy of Time*. Oxford: OUP.
- Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Wiley-Blackwell.
- Burns, Z. C., Caruso, E. M., & Bartels, D. M. (2012). Predicting premeditation: Future behavior is seen as more intentional than past behavior. *Journal of Experimental Psychology: General*, 141(2), 227.
- Byrne, C. A., Hyman, I. E., & Scott, K. L. (2001). Comparisons of memories for traumatic events and other experiences. *Applied Cognitive Psychology*, 15, 119–133.
- Caruso, E. (2010). “When the Future Feels Worse than the Past: A Temporal Inconsistency in Moral Judgement” *Journal of Experimental Psychology: General* (139):4 610–624.
- Caruso, E., Gilbert, D. T., and T. D. Wilson (2008). “A Wrinkle in Time: Asymmetric Valuation of Past and Future Events”. *Psychological Science* 19(8): 796-801
- D’Argembeau, A., Comblain, C., & Van der Linden, M. (2003). Phenomenal characteristics of autobiographical memories for positive, negative, and neutral events. *Applied Cognitive Psychology*, 17, 281–294.
- Destun, L. M., & Kuiper, N. A. (1999). Phenomenal characteristics associated with real and imagined events: The effects of event valence and absorption. *Applied Cognitive Psychology*, 13, 175–186.
- Dufwenberg M, Kirchsteiger G. 2004 A theory of sequential reciprocity. *Games Econ. Behav.* 47, 268–298. (doi:10.1016/j.geb.2003.06.003)
- Elster, J. (1986). *Rational Choice*, New York: NYU Press.
- Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7, 117–140.
- Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2), 351-401.

- Green, L., Fisher, E. B., Perlow, S., and Sherman, L. (1981). 'Preference reversal and self control: Choice as a function of reward amount and delay'. *Behaviour Analysis Letters*, 1, 43–51.
- Greene, P. (ms). 'Time Bias, Social Bias, and Future Generations.'
- Greene, P. and M Sullivan, (2015). 'Against Time Bias' *Ethics* (125)5: 947-970.
- Greene, P., Latham, A. J., Miller, K, and Norton, J. (ms). 'Hedonic and Non-hedonic Bias Toward the Future'.
- Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics*, 33-54.
- Hoch, S. J., & Loewenstein, G. F. (1991). Time-inconsistent preferences and consumer self-control. *Journal of consumer research*, 17(4), 492-507.
- Johnson, M. W., & Bickel, W. K. (2002). "Within-subject comparison of real and hypothetical money rewards in delay discounting". *Journal of the Experimental Analysis of Behavior*, 77, 129–146.
- Khaw, M. W., Glimcher, P. W., & Louie, K. (2017). Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, 201715293.
- Lagorio, C. H., & Madden, G. J. (2005). "Delay discounting of real and hypothetical rewards III: Steady-state assessments, forced choice trials, and all real rewards." *Behavioural Processes*, 69,173–187.
- Lewis, C. I. (1946). *An Analysis of Knowledge and Valuation*. LaSalle, IL: Open Court.
- Loewenstein, G. (1996). 'Out of Control: Visceral Influences on Behavior'. *Organizational Behavior and Human Decision Processes*, 65(3), 272–292.
- Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal*, 97(387), 666-684.
- Loewenstein, G., & Prelec, D. (1991). Negative time preference. *The American Economic Review*, 81(2), 347-352.
- Louie K, Khaw MW, Glimcher PW (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci* **110**:6139–6144
- MacKeigan, L. D., Larson, L. N., Draugalis, J. R., Bootman, J. L., & Burns, L. R. (1993). Time preference for health gains versus health losses. *Pharmacoeconomics*, 3(5), 374-386.
- Marsh, H. W. & Parker, J. W. (1984). Determinants of student self-concept: Is it better to be a relatively large fish in a small pond even if you don't learn to swim as well? *Journal of Personality and Social Psychology*, 47, 213-231.
- Martin, J. (1986). The tolerance of injustice. In Olsen J. M., Herman C. P. & Zanna M. P. (Eds.), *Relative Deprivation and Social Comparison: The Ontario Symposium* (pg. 217-242). Hillsdale, NJ: Erlbaum.
- Nagel, T. (1970). *The Possibility of Altruism*, Oxford: Clarendon.
- Nowak MA, Sigmund K. 1998 The dynamics of indirect reciprocity. *J. Theor. Biol.* 194, 561–574. (doi:10.1006/jtbi.1998.0775)
- Nowak MA, Sigmund K. 2005 Evolution of indirect reciprocity. *Nature* 437, 1291–1298.(doi:10.1038/nature04131)
- Pleban, R. & Tesser, A. (1981). The effects of relevance and quality of another's performance on interpersonal closeness. *Social Psychology Quarterly*, 44, 278- 285.
- Rabin M. 1993 Incorporating fairness into game theory and economics. *Am. Econ. Rev.* 83, 1281–1302.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Belknap.
- Read, D., & Van Leeuwen, B. (1998). Predicting hunger: The effects of appetite and delay on choice. *Organizational behavior and human decision processes*, 76(2), 189-205.

- Salovey, P. & Rodin, J. (1984). Some antecedents and consequences of social-comparison jealousy. *Journal of Personality and Social Psychology*, 47, 780-792.
- Sidgwick, H. (1884). *The Methods of Ethics*. London: Macmillan.
- Slote, A. (1983). *Goods and Virtues*. OUP.
- Soman, D., Ainslie, G., Frederick, S., Li, X., Lynch, J., Moreau, P., Mitchell, A., Read, D., Sawyer, A., Trope, Y., Wertenbroch, K., & Zauberman, G. (2005). "The Psychology of Intertemporal Discounting: Why Are Distant Events Valued Differently from Proximal Ones?", *Marketing Letters*, 16(3/4), 347–360.
- Suchman, Edward A. (1950), "The Intensity Component in Attitude and Opinion Research," in *Measurement and Prediction*, Samuel Stoffer et al., eds., Princeton: Princeton University Press.
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics letters*, 8(3), 201-207.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Velleman, J. D. (1991). "Well-Being and Time." *Pacific Philosophical Quarterly* 72 (1):48-77.
- Wills, T. A. (1981). Downward comparison principles in social psychology. *Psychological Bulletin*, 90, 245-271.
- Wark B, Lundstrom BN, Fairhall A (2007) Sensory adaptation. *Curr Opin Neurobiol* 17: 423–429.
- Wood, J. V., Taylor, S. E. & Lichtman, R. R. (1985). Social comparison in adjustment to breast cancer. *Journal of Personality and Social Psychology*, 49, 1169-1183.
- Yi, R., Gatchalian, K. M., and Bickel, W. K., (2006). "Discounting of Past Outcomes". *Experimental and Clinical Pharmacology*. 14(3): 311-317.