Against Time Bias

Greene, Preston; Sullivan, Meghan

2015

Greene, P., & Sullivan, M. (2015). Against Time Bias. *Ethics*, 125(4), 947-970.

<https://hdl.handle.net/10356/79718>

<https://doi.org/10.1086/680910>

© 2015 The University of Chicago. This paper was published in *Ethics* and is made available as an electronic reprint (preprint) with permission of The University of Chicago. The published version is available at: [<http://dx.doi.org/10.1086/680910>]. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper is prohibited and is subject to penalties under law.

Downloaded on 26 Apr 2023 19:14:49 SGT

Against Time Bias

Author(s): Preston Greene and Meghan Sullivan

Source: *Ethics*, Vol. 125, No. 4 (July 2015), pp. 947-970

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/10.1086/680910>

Accessed: 02/09/2015 04:36

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*.

<http://www.jstor.org>

Against Time Bias*

Preston Greene and Meghan Sullivan

Most of us display a *bias toward the near*: we prefer pleasurable experiences to be in our near future and painful experiences to be in our distant future. We also display a *bias toward the future*: we prefer pleasurable experiences to be in our future and painful experiences to be in our past. While philosophers have tended to think that near bias is a rational defect, almost no one finds future bias objectionable. In this essay, we argue that this hybrid position is untenable. We conclude that those who reject near bias should instead endorse complete temporal neutrality.

I. NEAR AND FUTURE BIAS

Good news—you've won a radio contest! There are two possible prizes, and the announcer is about to make the big reveal. One prize is a pleasant skiing vacation in Vermont to take place this weekend. The other prize is a lavish, week-long ski trip in the Swiss Alps five years from now. Suppose that you don't have any anticipated scheduling conflicts and that your prize—whichever it is—is assured and nontransferable. Which do you prefer? Many of us can handle some delayed gratification, but some may think five years is just too long to wait for a good thing.

Human beings are naturally disposed to discount the distant future, sometimes preferring worse experiences to better ones merely because

* Both of us contributed equally to this article, and we were equally involved in every stage of its conception and writing. We are grateful to Ruth Chang, Tom Dougherty, Andy Egan, Caspar Hare, Geoff Lee, Ben Levinstein, Callie Phillips, Kathryn Pogin, Mary Salvaggio, Meg Schmitt, Jeff Speaks, two referees, and the editorial board at *Ethics* for useful comments and guidance. The article also benefitted from substantial discussion with audiences at Northern Illinois University, Rutgers University, NTU Singapore, Oxford University, Leeds University, Western Washington University, the Massachusetts Institute of Technology, the Summer Institute for Diversity in Philosophy, the 2014 Inland Northwest Philosophy Conference, the Northwest Workshop on Time and Rationality, the University of Waikato, and the University of Sydney.

Ethics 125 (July 2015): 947–970

© 2015 by The University of Chicago. All rights reserved. 0014-1704/2015/12504-0002\$10.00

of their temporal proximity. How we discount changes as we grow and mature. For instance, most children discount the future at extraordinary rates. When presented with the choice of eating one marshmallow now or waiting a few minutes and receiving another, it can be quite a struggle for a four-year-old to delay gratification. Adults do a better job at short-term delayed gratification, but we still discount over longer periods of time—we fail to save enough for retirement, persist in smoking, and neglect to get minor dental problems corrected before they require root canals.

Future discounters display a *bias toward the near*. Stated in general terms:

An agent *S* is biased toward the near with respect to pleasure iff for two exclusive future experiences, E_1 and E_2 , where E_2 is at least as pleasurable as E_1 (adjusted for the subjective probabilities of the events occurring), *S* prefers E_1 because it would occur nearer to the present.

For pains the preference is inverted:

An agent *S* is biased toward the near with respect to pain iff for two exclusive future experiences, E_1 and E_2 , where E_1 is at most as painful as E_2 (adjusted for subjective probabilities), *S* prefers E_2 because it would occur further in the future.

We add the qualifications about probabilities because distant experiences are typically less certain than near ones. Future discounters are biased toward the near even after the probabilities have been taken into account.¹

Now consider a variant on the example.² Again, there is good news—you have won another radio contest! But this time there is also bad news—you suffer from temporary retrograde amnesia. As a result, you are not able to form short-term memories stretching more than a few days back. You know that you've either won and will enjoy the Vermont vacation

1. In addition to probabilities, there are several other factors that may have an effect on temporal preferences but do not essentially concern nearness. Here are two examples. First, an agent might prefer near pleasures and distant pains for reasons concerning personal identity (e.g., the agent might be skeptical that she will be the same person that feels distant pain and pleasure). Second, an agent might prefer that pleasures and pains occupy certain temporal locations because she has "global" preferences about the structure of her life (e.g., she prefers an improving life to a degrading one, regardless of the total amount of pleasure in each life). We remain neutral regarding the rationality of these types of preferences.

2. Inspired by the amnesia cases originally presented in Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984), pt. 2.

next month, or you've won and already enjoyed the Swiss vacation last month. Which situation would you prefer to be in?

Many of us strongly prefer the future-but-moderate vacation to the spectacular-but-past one. This reveals a *bias toward the future*. Stated in general terms:

An agent *S* is biased toward the future with respect to pleasure iff for two exclusive experiences, E_1 and E_2 , where E_1 is at least as pleasurable as E_2 , *S* prefers E_2 because it is a present or future pleasure rather than a past one.

And similarly for pains:

An agent *S* is biased toward the future with respect to pain iff for two exclusive experiences, E_1 and E_2 , where E_2 is at most as painful as E_1 , *S* prefers E_1 because it is a past pain rather than a present or future one.

Do we need to adjust for the subjective probabilities of the experiences occurring (or having occurred), as we did with the bias toward the near? This depends on how future biased one is. Many are tempted to assign no value to past experiences when reflecting on their preferences. For such agents, the probabilities do not matter.³

Are time-biased preferences rational? And is there a difference in the rationality of near and future bias, or should they be treated similarly? We will argue that they should be treated similarly and that they are both irrational. In the next section, we describe what we take to be the strongest case against the rationality of near bias. In the rest of the essay, we show that the same case can be made against future bias.

II. TWO METHODS FOR EVALUATING NEAR BIAS

Most of the literature on time bias to date has focused on the rational status of near bias. There are two competing views on the rationality of near bias: (i) the *economic view*, which focuses on dynamic consistency, and (ii) the *philosophical view*, which focuses on the practical costs and intuitive arbitrariness of near bias. Consider each in turn.

Standard theories of rational preference in decision theory and economics are primarily concerned with structural constraints on one's preferences. For example, it is thought that consistency requires that one's preferences be transitively ordered—if one prefers *A* to *B*, and *B* to *C*, then one also prefers *A* to *C*—and that they be complete, in the sense that

3. We discuss different ways one might discount the past in Sec. IV.

an agent has a preference or is indifferent between any two outcomes. It is a matter of some controversy which particular constraints should be adopted, but what characterizes the economic view is a laissez-faire attitude to nonstructural features (the content) of preferences. An agent can prefer, for example, the destruction of the world to the prick of her finger as long as all of her preferences cohere. This laissez-faire approach applies to the bias toward the near: as long as one's preferences are consistent, it is permissible to prefer the worse vacation because it is nearer.

Proponents of the economic view will, nevertheless, criticize near-biased agents who discount the future in dynamically inconsistent ways. A discount rate is a function for determining how much an agent discounts the value of an experience based on its temporal distance. Some discount functions (e.g., hyperbolic functions) can result in dynamic inconsistency: an agent might prefer a future experience *A* to another future experience *B* when both experiences are relatively far off but then reverse this preference as the experiences become nearer. Other discount functions (e.g., exponential functions) never lead to dynamic inconsistency and so are deemed rationally permissible on the economic view.⁴

In contrast to the economic view, many philosophers have concluded that near bias, even in its dynamically consistent forms, is rationally impermissible. One common complaint about near-biased preferences is that they are objectionably arbitrary. For example, Sidgwick in his discussion of the various formulations of his impartiality principle urges: "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)."⁵ Rawls reiterates the point in *A Theory of Justice*: "Rationality requires an impartial concern for all parts of our life. The mere difference of location in time, of something's being earlier or later, is not a rational ground for having more or less regard for it." In line with Sidgwick's caveat, he continues: "Of course, a present or near future advantage may be counted more heavily on account of its greater certainty or probability, and we should

4. For a philosophically accessible overview of this debate, see George Loewenstein and Jon Elster, eds., *Choice over Time* (New York: Russell Sage Foundation, 1992). 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 calls "pure" discounting—and the discounting of commodities used by economists in cost-benefit analysis. See John Broome, *Weighing Goods: Equality, Uncertainty, and Time* (Oxford: Blackwell, 1991). We mean to be exclusively discussing the rationality of pure discounting.

5. Henry Sidgwick, *The Methods of Ethics* (London: Macmillan, 1884), bk. 3, chap. 13, pp. 380–81.

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.”⁶ Distant experiences are typically less certain than near ones, and this difference in their probabilities may give us a reason to prefer near pleasures and distant pains over distant pleasures and near pains. However, there is no reason to prefer certain experiences based solely on their temporal position.⁷

By itself, the complaint that near bias is arbitrary seems to cover too much ground. For example, some have claimed that the arbitrariness considerations raised in favor of temporal neutrality equally support agent neutrality.⁸ This would be unfortunate for critics of near bias, since it is highly controversial whether rationality requires agent neutrality. So critics of near bias are presented with a challenge: why is the argument in favor of temporal neutrality more plausible than one in favor of agent neutrality?

We think the best response to this challenge invokes what David Brink calls “compensation.”⁹ Temporal neutrality is characteristic of prudence; specifically, in the farsighted trade of smaller present goods for greater future goods when the probabilities are right. The allure of prudence is not that prudent people steer clear of making arbitrary distinctions. Rather, it is that prudent people make choices that result in their leading better lives (featuring, *inter alia*, cushy retirements, better health, and extra marshmallows). For this reason we think that children ought to learn prudence—perhaps by better understanding the rewards to be had from waiting for treats and by learning psychological tricks for resisting temptation. (On the economic view, implausibly, children are simply agents who discount the future at steeper rates than the rest of us.) We also tend to criticize adults as being imprudent when they engage in longer-scale discounting. The crucial point is that when an agent makes prudent sacrifices, that same agent is usually compensated for the loss later. But agents who sacrifice their own smaller rewards to provide larger

6. John Rawls, *A Theory of Justice* (Cambridge, MA: Belknap, 1971), 293–94.

7. Other examples of philosophers who reject near bias include Broome in *Weighing Goods*; C. I. Lewis in *An Analysis of Knowledge and Valuation* (LaSalle, IL: Open Court, 1946); Thomas Nagel in *The Possibility of Altruism* (Oxford: Clarendon, 1970); and Jon Elster in his introduction to *Rational Choice* (New York: NYU Press, 1986).

8. In *The Possibility of Altruism*, Nagel famously relies on the apparent parity of the considerations for temporal neutrality and agent neutrality to argue against the rationality of agent bias. In *Reasons and Persons*, Parfit stops short of endorsing Nagel’s conclusion, but he does go to great lengths to show the difficulty in finding a coherent justification for temporal neutrality that does not also apply to agent neutrality.

9. David O. Brink, “Prospects for Temporal Neutrality,” in *The Oxford Handbook of Time*, ed. Craig Callender (Oxford: Clarendon, 2010), 353–81.

rewards for others do not (in many instances) have any chance of compensation; the trade-offs demanded by agent neutrality are not investments but rather pure costs. This connection between practical rationality and compensation thus gives us a way to separate arguments for temporal neutrality from arguments for agent neutrality.

To summarize: there are two important components to the philosophical rejection of near bias. First is the assumption that near-biased agents make arbitrary distinctions among future experiences. Second is the observation that such agents end up living worse lives because of their arbitrariness.

Does the same argument apply against future bias? Presumably, many philosophers believe that these criticisms do not apply to future bias. Indeed, the rationality of future bias is assumed in prominent arguments for why we should view death as bad, for why we should reject certain metaphysical theories of time, and for why it is acceptable to be self-interested.¹⁰ As Christopher Heathwood puts it, a future-biased agent “is being completely reasonable in preferring that his pain be in the past. In fact, even his no longer caring at all that it occurred is perfectly fitting—not at all inappropriate. Why should he care about it now? No reason—it’s over and done with.”¹¹

What could justify such radically different approaches to these two forms of time bias? One idea might be that the future and the past are different from each other in a way that the near and distant future are not, and so future bias is less arbitrary than near bias. This suggestion, however, is a matter of significant metaphysical dispute. It is true that according to a natural formulation of the A-theory of time, events can be distinguished based on intrinsic temporal properties; namely, whether they are past, present, or future. However, according to the most com-

10. Anthony Brueckner and John Martin Fischer appeal to future bias to argue that it is reasonable to have different attitudes toward death and prenatal nonexistence in “Why Is Death Bad?,” *Philosophical Studies* 50 (1986): 213–21. Arthur Prior appeals to future bias to argue that B-theories of time—which deny an objective past/present/future distinction—lack the resources to explain the rationality of attitudes like relief in his “Thank Goodness That’s Over,” *Philosophy* 34 (1959): 12–17. Caspar Hare uses analogies to future bias to argue for the reasonableness of self bias in “Self-Bias, Time-Bias, and the Metaphysics of Self and Time,” *Journal of Philosophy* 104 (2007): 350–73. Christopher Heathwood argues that fitting-attitude theories of welfare are refuted by the fact that they do not account for the rationality of future bias in “Fitting Attitudes and Welfare,” in *Oxford Studies in Metaethics*, ed. Russ Shafer-Landau, vol. 3 (Oxford: Oxford University Press, 2008), 47–73. Other philosophers like G. E. Moore and Lars Bergstrom build future bias into normative decision theories. See G. E. Moore, “A Reply to My Critics,” in *The Philosophy of G. E. Moore*, ed. Paul Arthur Schlipp (Chicago: Open Court, 1942), 535–677; and Lars Bergstrom, *The Alternatives and Consequences of Actions* (Stockholm: Almqvist & Wiksell, 1996).

11. Heathwood, “Fitting Attitudes and Welfare,” 56–57.

mon A-theory—presentism—only the present time and present objects are real. So, according to presentism, the past and future are metaphysically very similar in that neither exists. Presentism does not give us a metaphysical reason to favor future experiences over past ones.¹²

The metaphysical justification for future bias is even less clear if we adopt the B-theory of time. The B-theory holds that there is no way to intrinsically distinguish events as past, present, or future. Rather, the most we can say is that some events are earlier or later relative to others. B-theorists think of other times in the way we ordinarily think of other regions of space. Just as there is no special “here” in space, B-theorists deny there is a special “now” in time. So given the B-theory of time, distinguishing between past and future experiences can seem just as arbitrary as distinguishing between experiences that happen here and experiences that happen there.

In any event, it is unlikely that those who endorse future bias do so because of considerations regarding the metaphysics of time. A more promising defense of future bias might be that it does not lead to the sort of imprudent trade-offs associated with near bias. After all, while we can trade between the near and the distant future, we cannot trade between the future and the past. For this reason, it might be thought that future bias is practically irrelevant.

This assumption is mistaken. As we’ll argue, future bias can affect practical reasoning, primarily through its interaction with other principles of rational planning related to regret. Before turning to this, however, we will survey a recent attempt to show the negative practical consequences of future bias through its interaction with risk aversion. While we think this argument is problematic, it points the way to more general problems for future bias.

III. FUTURE BIAS AND RISK AVERSION

In an interesting 2011 essay in this journal, Tom Dougherty argues that future-biased agents who are risk averse can be turned into “pain pumps”: that is, they would accept a series of trades that guarantees they will suffer more pain overall and be better off in no respect. Since risk aversion is often taken to be rationally permissible, if Dougherty’s argument succeeds this would be a major challenge to the rationality of future bias. However, we will see that Dougherty’s argument is problematic because it requires that an agent get increasing marginal disutility from pain. Before explaining this in more detail, let us first consider Dougherty’s puzzle.

12. For additional discussion of how presentism relates to future bias, see Hare, “Self-Bias, Time-Bias, and the Metaphysics of Self and Time,” 350–73.

According to Dougherty, risk-averse agents are always willing to sacrifice some amount of expected value in order to reduce the amount of risk they face. Dougherty captures this idea with the following principle:

Every Risk Reduction Has Its Price: If a robustly risk-averse person faces a fifty-fifty gamble, then for any reduction of the gap between the good and bad outcomes of the gamble, there is some decrease in the gamble's expected value that this person would accept in return for this reduction of the gap.¹³

To illustrate this, Dougherty explains that a risk-averse agent would prefer to exchange a ticket that has an equal chance of paying \$20 or \$0 for a ticket that pays \$10 with certainty. In this case, such an agent would have received a reduction in the gap between the good and bad outcomes for free—both tickets have an expected value of \$10. However, a risk-averse agent might also be willing to exchange the first ticket for one of lower expected value: for example, a ticket with an equal chance of paying \$8.50 or \$10.50. Such a risk-averse agent would thus accept a \$.50 reduction in the expected value of their ticket in order to reduce the gap between the best and worst outcomes.

With this characterization of risk aversion in place, Dougherty presents a puzzle for agents who are risk averse and future biased. Here is a version of Dougherty's case:

Blue Pills and Red Pills: You know that you are facing one of two equally probable situations. If you are in Situation *A*, you will experience four hours of pain on Tuesday and two hours of pain on Thursday. If you are in Situation *B*, you will just experience three hours of pain on Thursday. You do not know which situation you are in, though you always know what day it is and you always know you will be offered the following choices.

On Monday you will be offered a blue pill with the following properties: if you take the pill and are in *A*, it reduces the time you suffer on Thursday by twenty-nine minutes. If you are in *B*, it increases the length of your pain on Thursday by thirty-one minutes. You are risk averse, so you will take the pill, intending to decrease the overall potential suffering of the worst situation: Situation *A*.

On Wednesday, you will have amnesia, and you will still be unsure what situation you are in. You will be offered a red pill with the following properties: if you take the pill and are in *A*, it increases the length of pain you suffer on Thursday by thirty minutes. If you are in *B*, it decreases the length of pain you suffer on Thursday by thirty minutes. You are future biased, and so from your perspective the

13. Tom Dougherty, "On Whether to Prefer Pain to Pass," *Ethics* 121 (2011): 525.

only pains that would be relevant are Thursday's. And you are risk averse, so you want the worst Thursday scenario—which is now *B*—to be improved. So you will also take the red pill.¹⁴

Because you will take both pills, you are guaranteed to suffer an extra minute of pain. If you are in Situation *A*, you will have decreased your suffering by twenty-nine minutes (blue pill) but then increased it by thirty minutes (red pill). If you are in Situation *B*, you will have increased your suffering by thirty-one minutes (blue pill) but then only decreased it by thirty minutes (red pill). Your preferences turn you into a pain pump.

Dougherty claims that his formulation of the puzzle is meant to be neutral between different decision-theoretic characterizations of risk aversion, but—given certain plausible assumptions—it is actually incompatible with the most common characterization. According to standard rational choice theory, rational risk aversion in relation to some good requires the good to have diminishing marginal utility or increasing marginal disutility. Monetary payoffs, it is thought, do indeed have diminishing marginal utility for most people—the utility of each additional dollar decreases as the total amount of a payoff increases. This is represented by a concave utility function, as seen in figure 1. This phenomenon would explain the cases concerning money with which Dougherty starts. Why would a risk-averse agent choose to exchange a fifty-fifty gamble between \$20 and \$0 for one that pays \$10 for certain? The answer, according to the standard view, lies in the diminishing marginal utility of money: the agent takes the utility of \$20 to be less than twice that of \$10. Therefore, according to the most common approach to risk aversion, the rationality (and the ubiquitousness) of risk aversion in relation to money is explained as a rational response to money's diminishing marginal utility.

But Dougherty's patient is trading in pain, rather than money. For standard decision theory to justify risk aversion in this situation, pain must have increasing marginal disutility—the disutility of each additional unit of pain must increase as the total amount of pain is increased. This is the only way, according to the standard approach, for it to be rational for an agent to agree to an increase in expected pain in exchange for taking the pills. For example, in order for it to be rational to take the blue pill, the agent would need to regard three hours of pain as having less than three-fifths the disutility of five hours of pain (keeping in mind that Dougherty assumes that the amount of pain provided by each hour is exactly the same).

Is this rational? It is hard to understand why an agent would have this attitude. While the effects of a painful stimulus might intensify as it

14. For Dougherty's original presentation of the case, see *ibid.*, 526–28.

FIG. 1

is applied over a period of time, an increase in actual units of pain provides one with exactly that amount of experienced pain. Therefore, we suggest that for any size units, two units of actual experienced pain is twice as bad as one unit of actual experienced pain, all else equal. And if so, then standard decision theory offers a ready-made response to Dougherty's puzzle: the irrationality lies with risk aversion to pain, and not future bias.¹⁵

It is possible that risk aversion to pain could be motivated without appeal to increasing marginal disutility. This would require a nonstandard framework. One example of an attempt to provide such a framework—which does seem compatible with Dougherty's arguments—is developed by Lara Buchak.¹⁶ Absent such a framework, Dougherty's argument fails to show that there is something irrational in future bias. Nevertheless, Dougherty's focus on preference instability points the way to more general problems with future bias and rational planning—problems that are unrelated to risk aversion.

15. We might call risk aversion in relation to experiences like pain and pleasure “pure” risk aversion. This is the analogue of pure discounting in the case of near bias. We believe that neither pure discounting nor pure risk aversion is permissible.

16. Lara Buchak, *Risk and Rationality* (Oxford: Oxford University Press, 2013).

IV. FUTURE BIAS AND REGRET

Anticipated regret can play a central role in how we plan for the future. Thinking about regret can be important in making far-reaching decisions—such as those regarding one’s career or marriage partner—but it also comes into play in day-to-day activities. Consider, for example, the experience of being tempted by something. Oftentimes temptations are hardly worth considering, but sometimes they are strong enough to induce a temporary reversal in preference. Michael Bratman provides an influential example of such a case:

The Second Pilsner: [Ann] enjoys a good read after dinner but also loves fine beer at dinner. However, she knows that if she has more than one beer at dinner she cannot concentrate on her book after dinner. Prior to dinner Ann prefers an evening of one beer plus a good book to an evening with more than one beer but no book. Her problem, though, is that each evening at dinner, having drunk her first Pilsner Urquell, she finds herself tempted by the thought of a second: For a short period of time she prefers a second beer to her after-dinner read. This new preference is not experienced by her as compulsive. If asked, she would say that right now she really prefers to go ahead this one time and have the second drink, though she will also acknowledge that even now she prefers that she resist similar temptations on future nights. As she knows all along, this change in ranking will be short-lived: after dinner she will return to her preference for a good read.¹⁷

Is it permissible for Ann to ignore the fleeting temptation? Or must she instead act on her preference for a second beer? It would be strange to criticize Ann for resisting the temptation to drink the second beer. After all, it seems like this behavior is characteristic of the sort of prudence that can have a positive impact on one’s life. In contrast, always acting on temptation will doom Ann to an unfortunate pattern of behavior: she will drink too much during dinner night after night, thereby never enjoying her book.

In the literature on regret and practical reasoning, avoiding regret is often viewed as the result of taking a more sophisticated perspective on one’s preferences through time. In this context, “regret” is understood as a type of preference—namely, preferring that one had done

17. Michael Bratman, “Toxin, Temptation, and the Stability of Intention,” in *Faces of Intention* (Cambridge: Cambridge University Press, 1999), 74. For an influential earlier discussion of some of the issues raised in this section, see Jon Elster, *Ulysses and the Sirens* (Cambridge: Cambridge University Press, 1979), 36–111.

otherwise—rather than as an affective or cognitive state.¹⁸ By taking a wider perspective on her preferences, and thereby noticing how her preferences are likely to shift through time, Ann is able to avoid the problems caused by preference change. Some theorists, like Bratman, take regret avoidance in cases like Ann's to be rationally required.¹⁹ A no-regrets requirement, according to Bratman, is most plausible when restricted to "no-unanticipated-information cases."²⁰ Thus, the rationality of avoiding regret is most obvious when agents are certain of the outcomes of the various options available to them.

For our purposes, we are not concerned with rational requirements, but rather only with the rational permissibility of regret avoidance. We propose the following principle:

Weak No Regrets: If an agent has full and accurate information about the effects of the options available to her, then it is rationally permissible for her to avoid options she knows she will regret in favor of ones she knows she will never regret.

Consider what it would mean to deny this claim. One would have to think that there is some sort of rational requirement against avoiding certain regret, even when agents have full information. (Perhaps one might think that rational agents are required to consider only the preferences of the present.) It is difficult to see the motivation for such a requirement, especially in light of the practical problems it would create.²¹

18. In common usage, "regret" is used to refer to many things, including a type of painful affective state, or a cognitive state involving the judgment that one's prior decision was irrational. However, the connection between Ann's preference structure and the generation of these states is, at best, contingent.

19. Frank Arntzenius endorses the principle "A rational person should not be able to foresee that she will regret her decisions"; see his "No Regrets; or, Edith Piaf Revamps Decision Theory," *Erkenntnis* 68 (2008): 277–97. See also Nagel, *Possibility of Altruism*, 57–76; Rawls, *Theory of Justice*, 416–23; and Graham Loomes and Robert Sugden, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal* 92 (1982): 805–24. Technically, Bratman holds that regret avoidance is required if and only if we assume, plausibly, that before dinner Ann intends to resist a second beer.

20. Bratman, "Toxin, Temptation, and the Stability of Intention," 79.

21. Perhaps the only examples that have the potential to give weak no regrets trouble are situations in which an agent anticipates experiencing irrational regret. For example, what if an agent is faced with the choice between killing herself or taking a "regret pill," which causes those who take it to regret doing so? Is the agent rationally required to take the pill? We don't think this is clear-cut. However, if the agent is so required, then weak no regrets seems to give the wrong result. In light of the possibility of this sort of case, we might consider further weakening weak no regrets to only apply to cases that do not feature anticipated irrationality. In any event, all of the examples we discuss in this essay do not feature anticipated irrationality.

A. *The Scheduling Problem*

Why is thinking about regret relevant to the rationality of future bias? For future-biased agents, reversals in preference are caused solely by the passing of time, and this generates cases in which they exhibit strange patterns of regret. Consider the following case:

Fine Dining: Jack wins a free meal at a fancy French restaurant on Monday morning, and he must schedule the meal for a night sometime in the next week. Given his flexible schedule, every night is equally convenient for him, and there are no other considerations that would make the meal more enjoyable or more likely to occur on one night rather than another. Therefore, Jack schedules the meal for Monday night. As expected, it is an incredibly delicious meal. On Tuesday morning, Jack strongly prefers that his restaurant experience were in the future, rather than the past. And so he regrets scheduling the meal for the previous night.

Jack's regret is rational according to defenders of future bias. At the time of his decision each choice available to him (the seven days in which the meal could be scheduled) involves the prospect of a future pleasurable experience. But on Tuesday only six of the seven options feature a future meal. Therefore, Jack prefers that he had not scheduled the meal for Monday. In fact, unless Jack schedules the meal for the last day allowed by the restaurant, he is assured that he will come to regret his choice. Therefore, if Jack is future biased and aims to avoid regret, then he must schedule the meal for the last day possible.

Furthermore, in cases in which there is no deadline, it is not clear when Jack should schedule the meal. If we abstract away from concerns over Jack dying, losing his sense of taste, the restaurant closing, and so on, then it appears that attempting to avoid regret will lead Jack to wait a very long time to have his meal. Imagine that every day Jack is asked whether he would like to schedule the meal the following day. Since Jack knows that he will come to regret scheduling it, he chooses not to schedule it each time. If Jack is unlucky enough to live forever, he may find that he never has the meal. This creates a paradox similar to that faced by agents who are constantly offered positive rates of return on their savings.²² Call this the *scheduling problem for future bias*—future-biased

22. Tjalling Koopmans calls this "the paradox of the indefinitely postponed splurge." See his "Objectives, Constraints, and Outcomes in Optimal Growth Models," *Econometrica* 35 (1967): 1–15. The problem is also similar to Frank Arntzenius, Adam Elga, and John Hawthorne's example "Trumped," presented in "Bayesianism, Infinite Decisions, and Binding," *Mind* 113 (2004): 251–83. In *Fine Dining*, however, Jack's payoff does not increase through time, and so this behavior seems even more problematic.

agents who avoid regret will postpone positive experiences for no good reason.

It is worth noting that we can get to the scheduling problem from a related starting point.²³ Suppose you endorse the following principle governing rational planning:

Metapreference Principle: It is rationally permissible for an agent to prefer a plan in which all of her present and future preferences are satisfied over a plan in which some are not.²⁴

The principle should seem intuitive: there is no rational requirement against making sure that all of your present and future preferences are satisfied, rather than merely a subset. But in *Fine Dining*, if Jack has this metapreference and is future biased, then he will be compelled to schedule the meal as late as possible. Indeed, the maximally deferred meal is the only plan that could satisfy all of his preferences. Thus, future-biased agents who have this metapreference will schedule their experiences in an absurd way.²⁵

The scheduling problem has gone unnoticed, we believe, because it does not apply to agents who are both future biased and near biased, and most of the literature on time bias has been concerned with rejecting the rationality of near bias. Our discussion thus builds on the rejection of near bias, and—in line with the philosophical view—it concerns agents who are future biased but near neutral. This dynamic generates the problem.

B. *The Meager Returns Problem*

The problems with regret avoidance and future bias run deeper than the scheduling problem: this combination can also lead to large deductions

23. We would like to thank a referee for suggesting this alternative route.

24. Rachael Briggs offers a similar principle in the context of rational decision making as an instance of an agent looking out for her “future selves”: if all of the agent’s possible future selves like *A1* at least as well as *A2*, and some of the agent’s possible future selves prefer *A1* to *A2*, then the agent should prefer *A1* to *A2*. See her “Decision-Theoretic Paradoxes as Voting Paradoxes,” *Philosophical Review* 119 (2010): 1–30.

25. To be more cautious: assume that an agent has the metapreference and consider three possible answers to the question “how should the agent schedule pleasurable experiences when the total pleasure is guaranteed to be the same?” One answer is that the agent should schedule the experiences as soon as possible (which is supported by actual practice, assuming that most of us are near biased). Another answer is that the agent should be indifferent as to when the experiences are scheduled (which we believe is supported by reflection). The remaining answer is that the agent should schedule the experiences as late as she possibly can. It is not clear what could motivate this last view, other than the combination of future bias and the metapreference principle. Therefore we should be suspicious of at least one of these principles.

in agents' overall happiness. This is because such agents will be willing to trade something to have their bizarre scheduling preferences satisfied. Consider Billy, who is offered a choice between two cookies immediately or one cookie at some point in the future. It would seem that the rational choice is clear: Billy should choose to have two cookies now. But for future-biased agents the answer is not so simple.

Assume that Billy is absolutely future biased: an absolutely future-biased agent assigns no value whatsoever to past experiences.²⁶ If he chooses to have two cookies now, then as soon as the cookies are consumed he will prefer that he had chosen to have one cookie later, because that experience would still be in his future. When the time at which he would have eaten one cookie passes—so that it is also in the past—he will become indifferent to his choice. Therefore, Billy can expect to regret choosing more cookies, and never regret choosing fewer cookies. So the only course of action that Billy will not come to regret is choosing to wait for fewer cookies.

In fact, for an absolutely future-biased agent it does not matter how many more cookies are available earlier. As long as Billy is interested in the later cookies at all, he will be forced to choose them or end up regretting his choice. We can imagine a case in which Billy is choosing between having ten cookies now or having just a morsel of cookie later. Again, he will come to regret choosing to have the ten cookies, but he will not come to regret choosing to have the morsel.

Therefore, if they aim to avoid regret, agents with an absolute future bias will be willing to trade the most spectacular of immediate pleasurable experiences for extremely mundane ones in cases like Billy's. But this is irrational.²⁷ Call this the *meager returns problem for future bias*.

C. *Could Nonabsolute Future Bias Avoid These Problems?*

In an attempt to avoid these problems, we might try denying the rationality of absolute future bias but still permitting a *nonabsolute* future bias.

26. Such an agent may of course assign value to present or future memories of those experiences. For simplicity, we will assume that this factor is not relevant in Billy's case. It is also interesting to note that for absolutely future-biased agents, subjective probabilities of the experiences do not matter—any chance of a future pleasure will have more value than a past pleasure. The permissibility of absolute future bias has defenders, as we saw with Heathwood's "Fitting Attitudes and Welfare."

27. The discussion of Billy's case here and below focuses solely on how the temporal locations of potential cookie-eating experiences affect Billy's preferences. There may, of course, be other factors that have an effect on Billy's preferences. For example, Billy might also prefer to be the sort of person that has chosen to eat more cookies in the past. If this preference outweighs the effects of future bias then Billy may avoid choosing fewer cookies. However, there is no guarantee that such outweighing factors will always be present to mask the problems that future bias creates.

A nonabsolutely future-biased agent assigns some value to past pleasures and pains, though it is less than that assigned to future ones.²⁸ For example, a nonabsolute future-biased agent might discount all past pleasures by a flat rate of 50 percent in comparison to future ones.

This kind of flat-rate future bias won't avoid the scheduling or meager returns problems. Recall that in *Fine Dining*, Jack's experience will be equally pleasurable regardless of when he schedules it, and so as long as he discounts past pleasures at all, he will prefer that he had scheduled the meal further in the future as soon as it has passed. Switching to Billy's cookie case, we need only make a simple adjustment to the presentation to account for nonabsolute future bias—we need to change the difference in value between the two experiences. While an agent who discounts past experiences 50 percent may not regret choosing two cookies now over one cookie in the future, he will regret choosing three cookies now over two cookies in the future. Similar reasoning shows that regardless of the discount factor used, there will exist a case in which the agent comes to regret choosing more pleasure sooner.

There is another way that one might have a nonabsolute future bias—one might apply a discount rate to past experiences. After all, when we discount the more distant future we do not do it absolutely or by a fixed value, but instead by applying a discount rate. What would follow if we did the same for the past?

Regardless of the past discount function they use, past discounters cannot avoid the scheduling problem. To illustrate this, again consider *Fine Dining* and imagine that Jack is choosing between scheduling the meal on Tuesday and scheduling it on Thursday. Figure 2 shows how Jack's preferences for these two potential meals might evolve if he applies a past discount function. If Jack schedules the meal for Tuesday, then as soon as the meal becomes past it will start to be discounted. Given that the meal is equally valuable whether it occurs Tuesday or Thursday, Jack will immediately prefer that he had scheduled the meal for Thursday, and he will never prefer that he had scheduled it for Tuesday. Further applications of this argument show that Jack must schedule the meal on the last day possible.

Do future-biased agents who apply past discount functions avoid the meager returns problem? At least in Billy's case, it appears that they do. If Billy applies a past discount function, then after he consumes the larger number of cookies there will be an interval in which he prefers the past cookie experience to the future one. Due to this, no matter what

28. Again, however, the importance of past experiences must not be tied solely to the expectation of pleasant or unpleasant future memories of the experiences.

FIG. 2

Billy chooses, there will be future times in which he prefers more cookies in the past.²⁹

Even though this kind of past discounting avoids the meager returns problem in Billy's case, it is problematic for more fundamental reasons. To see why, consider one of the most celebrated motivations for future bias—Parfit's *My Past or Future Operations*:

I am in some hospital, to have some kind of surgery. Since this is completely safe, and always successful, I have no fears about the effects. The surgery may be brief, or it may instead take a long time. Because I have to co-operate with the surgeon, I cannot have anaesthetics. I have had this surgery once before, and I can remember

29. However it is possible to construct cases in which the meager returns problem applies to past discounters. Suppose that the pleasure that Billy gets out of a cookie is a function of the size of the cookie, and that the size of a cookie continuously diminishes over time. And suppose that Billy is offered the choice between eating one (intact) cookie now or waiting for a crumbling cookie. It seems obvious that Billy should prefer the intact cookie. But suppose Billy discounts the past at a rate of $f(t) = e^{-\alpha t}$. And suppose the cookie crumbles at the slower rate of $f(t) = e^{-\beta t}$. Then at every future interval, Billy will prefer to wait for the crumbling cookie. But this is absurd. Such a scenario could be constructed for any discount rate Billy employs, because for any downward sloping curve, it is possible to draw another curve that is downward sloping but less steep (the crumble curve). We are grateful to a referee for suggesting this possibility, as well as to Ben Levinstein and Mary Salvaggio for discussion.

how painful it is. Under the new policy, because the operation is so painful, patients are now afterwards made to forget it. Some drug removes their memories of the last few hours.

I have just woken up. I cannot remember going to sleep. I ask my nurse if it has been decided when my operation is to be, and how long it must take. She says that she knows the facts about both me and another patient, but that she cannot remember which facts apply to whom. She can tell me only that the following is true. I may be the patient who had his operation yesterday. In that case, my operation was the longest ever performed, lasting ten hours. I may instead be the patient who is to have a short operation later today. It is either true that I did suffer for ten hours, or true that I shall suffer for one hour.

I ask the nurse to find out which is true. While she is away, it is clear to me which I prefer to be true. If I learn that the first is true, I shall be greatly relieved.³⁰

The intuition here is clear: it is much better to have a more painful surgery in one's past rather than a less painful surgery in one's future. However, for exponential and hyperbolic past discounters, this judgment depends on how far in the past the more painful surgery is. This is revealed by figure 3. In this example, whether the past surgery is preferable to the future surgery depends on where the agent is temporally located, since at some point between the two surgeries the curves representing disvalue cross. Imagine that the point at which they cross is about five hours after the more painful surgery. Given this, whether the future surgery is preferable to the past surgery would depend on whether more or less than five hours have elapsed since the more painful surgery. If it has been four hours since the more painful surgery, then the agent prefers the less painful future surgery. However, if it has been six hours, then the agent prefers the more painful past surgery.

In the original example, it is supposed to strike us as absurd for Parfit to hope for a future surgery. We submit that it is even more absurd for Parfit's hopes to rest on how much time has elapsed since when the past surgery would have taken place.³¹ Past-discounting models of future bias are therefore implausible.

Let's summarize. The argument against future bias in this section takes the following form. (1) It is permissible to avoid certain regret. (2) If

30. Parfit, *Reasons and Persons*, 165.

31. Furthermore, the attitudes of agents who apply past discount rates appear bizarre even when there is no amnesia. Consider that in many cases, such agents would find the future surgery preferable to the past surgery as soon as the past surgery ends, but then slowly come to prefer the past surgery as it drifts further into the past. This is in direct opposition with the "thank goodness that's over" intuition often cited in support of future bias.

FIG. 3

one is future biased and chooses to avoid certain regret, then one acts in ways that lead to the scheduling and meager returns problems. (3) It is irrational to act in such ways. (4) Therefore, future bias is irrational.

Consider what it would mean to reject premise 1. One might argue that there is a rational requirement against avoiding certain regret. But this is implausible. The other option is ad hoc—one might hold that one can aim to avoid regret except in cases in which doing so leads to the scheduling or meager returns problems. Since both of these options are difficult to accept (and since the scheduling problem can also be generated using the metapreference principle), we tentatively conclude that future bias—not regret avoidance—is the culprit. In the next section, we strengthen this conclusion by showing that future bias is susceptible to a plausible theory of error.

V. EXPLAINING THE PREVALENCE OF NEAR AND FUTURE BIAS

We have shown some of the ways that near and future bias are subject to the same kind of rational criticism. In this final section we add one more. Any time a theory postulates widespread unnoticed irrationality it owes us a theory of error—in this case, a theory for why certain preferences seem rational to us even though they are not. We will argue that each form of time bias has a similar underlying explanation.

A. Near Bias as a Probability Heuristic

Philosophers and psychologists have favored the idea that near bias is generated by a probability heuristic. While it must be granted that temporally distant experiences are usually less certain, and that we rationally ought to discount future experiences based on the probability that they will occur, it would still be a mistake—as Sidgwick and Rawls stress—to hold that the preferability of the experiences themselves is dependent on their distance from the present. Rather, when we consider trading between experiences in our near and distant future, the two factors to consider are the value of the experiences and their probabilities. The problem is that probabilities are difficult to compute, and this creates the need for a heuristic. In order to deal with this problem, we apply a discount function to experiences based on their distance from the present, thus making decision situations tractable while creating an approximation of the probabilities of the experiences. While doing this typically leads us in the right direction, it nevertheless backfires in situations in which there is a significant divergence between probabilities and our rate of discounting.

There is also an evolutionary story to be told about how we developed the emotions associated with near bias. Emotions can be viewed as adaptations whose purpose is to solve basic ecological problems facing organisms.³² On the evolutionary hypothesis, near-biased anxiety and relief evolved as adaptations for tracking probabilities.³³ These emotions play different roles in our decision making at different stages of our lives. As we age, we may experience the emotions associated with near bias to a lesser degree, but they do not disappear—we must learn to overcome them. Psychologists studying emotional regulation identify the mechanisms by which we do this.

The foregoing highlights the reasons why emotional responses should not necessarily be understood as trustworthy guides to rational preference. Instead, rational agents ought to reflect on how their emotions are leading (or misleading) them when they form preferences. To better illustrate this, consider a common trigger of near-biased emotions: the dentist. Imagine that you are given a choice between undergoing a moderately painful dental surgery tomorrow or delaying the surgery for a year, risking the problem becoming worse. Prudently, you choose to make an appointment for tomorrow. However, you may find yourself thinking about tomorrow's surgery, and even feeling anxiety, and this

32. See Charles Darwin, *The Expression of the Emotions in Man and Animals* (London: Harper Collins, 1998); Robert Frank, *Passions within Reason: The Strategic Role of Emotions* (New York: Norton, 1998); Robert Plutchik, *Emotion: A Psychoevolutionary Synthesis* (New York: Harper & Row, 1980).

33. See Loewenstein and Elster, *Choice over Time*.

anxiety may build as the time of the appointment draws nearer. At the same time, you know that if you had scheduled the distant surgery, you would not be anxious. Does the fact that you would feel anxiety only about the near surgery show that the distant surgery is preferable to the near one? Of course not—your preference for the near surgery should remain.

Now imagine that right before the drill is about to go in, the dentist informs you that due to a complicating factor the surgery will need to be postponed for a year. You feel relieved. Does the fact that you feel relief show that this distant surgery is preferable to the near one? Again, the answer is no. Your relief, like your anxiety, is not a good guide to what your preferences should be. It merely reveals that your rational commitments sometimes have little effect on your emotional responses.

B. Future Bias as a Control Heuristic

When we switch from thinking about near bias to thinking about future bias, it is clear that we have the same sort of temporally biased emotions, but the significance of these emotions is understood very differently. In fact, for some reason the relief and anxiety associated with future bias is taken to be evidence for its rationality. For example, Parfit takes the main obstacle to rejecting future bias to be that it requires us to think that he would be “irrational to be relieved when suffering is in the past” when he learns that he has already had the surgery.³⁴ But this is equivalent to the situation involving near bias just described: it is easy to imagine someone being relieved to hear that his dental surgery needs to be postponed for a year, and we need not think that this person is irrational. Rather, we might think that his emotional response is hardwired in a way that is not under the control of his rational commitments. Why should we think otherwise in the case of future bias?

It might seem that the way one thinks about relief will have an effect on the plausibility of Parfit’s claim, since, just as with regret, there are many ways to understand what it means to be “relieved.” On one extreme we might think of relief simply as a kind of pleasant affective state, while on the other extreme we might think of it as a highly cognitive state involving the judgment that one is better off because the surgery is in the past. However, for our purposes it does not matter how one understands the relief featured in Parfit’s case, since it is sufficiently similar to the relief caused by the delay of the dental procedure. Regardless of how the relief caused by the delay is interpreted, it does not support the rationality of near bias. On the one hand, if relief is thought to involve judging that one is better off because the procedure must be delayed, then it is not clear that this is rational. On the other hand, if such relief is under-

34. Parfit, however, ultimately remains neutral regarding the rationality of future bias. See his *Reasons and Persons*, 186 and 194.

stood only as a “feeling,” then it seems appropriate to view this feeling not as rational or irrational, but rather as a nonrational response to one’s situation. Therefore, the presence of relief does not support the rationality of near bias, and neither should it support the rationality of future bias.

Furthermore, as with near bias, there are plausible evolutionary hypotheses about the emotions associated with future bias.³⁵ Just as it is advantageous for an agent to make trade-offs that respect the relative probabilities of potential rewards and penalties, it is also advantageous for an agent to focus more attention and energy on what is within her control. This evolutionary explanation can be thought to generate a control heuristic for future bias: future-biased emotions and preferences evolved to track asymmetries in control. A policy of not caring very much about the past is part of a good strategy for focusing on what is within our control, since future events are sometimes under our control while past events never are.

Evolutionary accounts of time bias are a relatively recent development, but the idea that future bias is really a form of control bias goes back at least to Hume. In the *Treatise*, he writes: “There is a phenomenon of a like nature with the foregoing, viz the superior effects of the same distance in futurity above that in the past. This difference with respect to the will is easily accounted for. As none of our actions can alter the past, ’tis not strange it shou’d never determine the will.”³⁶ And some of the most compelling evidence in favor of the control heuristic comes from the way our judgments change when we stop thinking about ourselves and instead consider the situation of other people. Note how our intuitions change when we consider distant people whose circumstances are beyond our control.³⁷ Suppose your friend from graduate school lives on the other side of the world. You hear through a mutual friend that he requires a painful surgery, which either occurred yesterday and was the longest ever recorded (ten hours) or will occur tomorrow and be much shorter (about an hour). Your informant cannot remember which of

35. James Maclaurin and Heather Dyke propose an evolutionary model for these emotions, but they do not use it to criticize the rationality of time bias; rather, it is part of their defense of the B-theory against Prior’s “Thank Goodness That’s Over” argument. See their “‘Thank Goodness That’s Over’: The Evolutionary Story,” *Ratio* 15 (2002): 276–92. The evolutionary model is also suggested in Paul Horwich, *Asymmetries in Time* (Cambridge, MA: MIT Press, 1992), 196–98; and Christopher Suhler and Craig Callender, “Thank Goodness That Argument Is Over: Explaining the Temporal Value Asymmetry,” *Philosophers’ Imprint* 12 (2012): 1–16.

36. David Hume, *A Treatise of Human Nature* (Oxford: Oxford University Press, 2000), sec. 2.3.7.6.

37. This was first noticed by Parfit, *Reasons and Persons*, 181–84. For further discussion of these kinds of cases, see Hare, “Self-Bias, Time-Bias, and the Metaphysics of Self and Time,” 350–73.

these is true. In this case, it is easy to imagine preferring that your friend did not have the longer surgery, even though it would at this point be in his past. Someone who thinks that future pain is worse than past pain might conclude that you are a rotten friend.

There is also strong empirical support for the presence of this self-regarding/other-regarding asymmetry in our future bias. Researchers at Harvard conducted a study in which they asked participants to determine fair compensation in different situations. First, some participants were asked to imagine that they had to do a boring data entry job one month in the future, and others were asked to imagine that they had completed the job a month ago. As expected, these participants exhibited future bias, believing they should be paid 60 percent more for future suffering. However, when participants were asked to compensate strangers for boring data entry work, they recommended the same compensation regardless of when the suffering occurred.³⁸

Self-regarding/other-regarding asymmetries are often taken as evidence of irrationality in the case of near bias. For example, consider Adam Smith's view of near bias:

In his steadily sacrificing the ease and enjoyment of the present moment for the probable expectation of the still greater ease and enjoyment of a more distant but more lasting period of time, the prudent man is always both supported and rewarded by the entire approbation of the impartial spectator, and of the representative of the impartial spectator, the man within the breast. The impartial spectator does not feel himself worn out by the present labour of those whose conduct he surveys; nor does he feel himself solicited by the importunate calls of their present appetites. To him their present, and what is likely to be their future situation, are very nearly in the same manner. He knows, however, that to the persons principally concerned, they are very different from being the same, and that they naturally affect *them* in a very different manner.³⁹

According to Smith, any inclination to be near biased disappears when we adopt the perspective of a spectator freed from the emotions that near pleasures and pains commonly elicit. We propose that the same is true of future bias: from an impartial perspective the inclination to be future biased evaporates. In keeping with Smith's style, we submit that specta-

38. See Eugene M. Caruso, Daniel T. Gilbert, and Timothy D. Wilson, "A Wrinkle in Time: Asymmetric Valuation of Past and Future Events," *Psychological Science* 19 (2008): 796–801. Suhler and Callender use this study to criticize Prior's "Thank Goodness That's Over" argument in their "Thank Goodness That Argument Is Over."

39. Adam Smith, *The Theory of Moral Sentiments* (Oxford: Oxford University Press, 1976), pt. 6.

tors are not “worn out” by the emotions that typically accompany future bias; that is, temporal differences in third-person cases do not create the same levels of anxiety or excitement as that experienced in first-person cases. Without these affect-laden mental states distorting our preferences we are inclined toward complete temporal neutrality.

VI. CONCLUSION

We have argued that rationality requires not just the rejection of near bias, but complete temporal neutrality. First, any kind of time bias (near or future) has the potential to cause agents to behave in irrational ways. Second, there is good reason to favor a unified approach to time biases, viewing both near and future bias as useful heuristics that nevertheless can lead us astray. We think there is much to recommend this outlook, and the burden now ought to be on defenders of future bias to show otherwise.