Is Discounting for Tense Rational?

Craig Callender

January 8, 2021

When we make decisions we are invariably comparing outcomes that happen at different times. How much should you sacrifice now to get a better job later? Should you switch to solar? Purchase a gym membership? Studies of intertemporal decision-making suggest that we often exhibit two types of time preferences: *future discounting*, that all else being equal, we prefer that future pleasures happen sooner than later (and vice versa for pains); and *past discounting*, that all else being equal, we prefer that pleasures happen in the present or future than in the past (and again, vice versa for pains).[1]

Are these time preferences rational? It's important that we make progress on this question, for assumptions about what discounting is normatively optimal inform public policy decisions throughout the world. Both social science and philosophy discuss the normative standing of discounting, philosophy focusing mostly on past discounting and social science mostly on future discounting. To a very rough first approximation, the two fields appear to disagree on when or if temporal discounting is rational. Future discounting is judged irrational by philosophers and as often rational by social scientists. Past discounting, by contrast, is viewed as rational by some philosophers but as (probably) irrational by social scientists.

Are the two fields really in disagreement? Part of the problem in determining if this is so is that they speak different languages about discounting. With the goal of bridging this divide, I want to focus on the fact that both time preferences are typically *tensed* in nature. This point is often obscured, so I think it's important to highlight it independently of the present interest. Through the lens of tense, we'll see that the conflict between social science and philosophy on discounting is mostly only apparent; and where genuine, it involves disputes within each field as much as between fields. Although I will not solve the question of whether discounting for tense is rational, keeping an eye on tense will allow us to make several novel observations, observations that I hope will help us better understand temporal discounting.

[1] Terminological confusion reigns over these two asymmetries. Many philosophers call future discounting "near bias" and past discounting "future bias". Since the topic of this paper is in part whether these preferences *are* biases, understood in a prejudicial way, I use more neutral terminology.

1 Background

Examples of future discounting include common preferences such as wanting to treat oneself to dinner tonight rather than wait until next week and preferring that the dentist trip happen next week instead of today. Social scientists have studied this kind of discounting in thousands of articles and experiments. The hope is to eventually intervene on many types of unfortunate individual and social outcomes, such as lack of saving, unhealthy eating habits, poor financial decisions, and addiction. In the paradigm that has developed, economists hold that future discounting can be rational, so long as it takes a special form, namely, exponential discounting (explained below). Psychologists and behavioral economics then point out that our actual discounting is typically suboptimal; they often try to find functions – often called "hyperbolic" even if not strictly hyperbolic – that fit the experimental data. Philosophers by contrast commonly dismiss future discounting as irrational, even obviously irrational.[2]

The situation is a bit the reverse with past discounting. Examples of past discounting include the familiar feeling that once an experience has happened, it is over and done, no longer an object of as much concern. Social scientists have not paid much attention to this type of discounting, but in philosophy it has received attention ever since Derek Parfit posed thought experiments that elicit in many people strong intuitions in favor of trading some small future pain for much greater past pain (Parfit 1984). Philosophers divide about whether past discounting is rational. So-called *temporal neutralists* (explained in Section 4) hold that one should not take temporal perspective into account when deciding the best location of goods and harms across a lifetime. Parfit himself is not entirely clear where he falls, but he is usually taken to be a non-neutralist, holding that past discounting is rational. Social scientists don't engage much with this kind of discounting because they think it is irrelevant to practical decisions. We'll see that if judged by the standard that they employ for future discounting then they would deem past discounting irrational.[3]

It seems like we have a clash between two traditions. Philosophers typically regard future discounting as irrational and put their energy into deciding whether past discounting is; by contrast, social scientists think a type of future discounting is rational and ignore past discounting.

[2]For a history of this literature, see Loewe 2006 and Peart 2000, and for some philosophical connections, see Żuradzki 2016. For a sample of its application, see Story et al 2014. Classic economic treatments of future discounting include Samuelson 1937, Ramsey 1928, Strotz 1955, and Koopmans 1960. For work in psychology challenging the standard economic model's descriptive accuracy, Thaler 1981 and Loewenstein and Prelec 1992 are classics, and Urminsky and Zauberman 2016 is an excellent review. There are scores of philosophers who hold that future discounting is irrational; for a sample see Ahmed 2018, Brink 2015, Bykvist 2015, Parfit 1984, Sullivan 2018, and against some of them, Dorsey 2017.

[3]The first empirical study of past discounting of which I'm aware is Caruso, E., Gilbert, D. and Wilson, T. 2008. As I write there are many empirical studies coming out in the "experimental philosophy" tradition that probe many of the claims often made about past discounting; see (e.g.) Lee et al 2020 and Greene et al 2020 for an entry into this literature. Lee et al 2020 report that the preference Parfit notices is not as strong as is often thought. For philosophical discussion, see Brink 2011, Bykvist 2015, Fernandes 2019, Hare 2015, Maclaurin and Dyke 2002, Parfit 1984 and Suhler and Callender 2012, as well as papers in this volume by Green et al and Lee and McCormack.

To make progress on understanding this apparent disagreement, I want to draw the reader's attention to the little-noticed fact that both types of time preference are typically *tensed* in nature. When we conceptualize time, we can characterize temporal relationships either in a tensed or tenseless manner. When we classify events in a tensed way, we identify whether they are past, present, or future; when we classify them in a tenseless manner, we instead place them in a sequence ordered by the the earlier than relation. Over a century ago McTaggart 1908 dubbed the first conceptualization using {past, present, future} an *A-series* and the second using {earlier than, later than, simultaneous with} a *B-series*. Meanwhile in semantics and cognitive linguistics the first is sometimes called *deictic time*, referring to its need for a deictic center, the Now, which is typically the time of utterance, and the second is often called *sequence time*. An easy way to tell the difference between the two conceptualizations is that tensed expressions change truth value depending on when they are said, whereas tenseless expressions do not. For instance, "I will see comet Neowise" was true only before I saw it but false afterward, whereas "the assassination of Lincoln is before that of Kennedy" is always true.[4]

Past discounting is manifestly tensed. When an event is over and done, what has happened is that the event *has become past*. It is not merely that the event is earlier than another. In Prior's famous 1959 discussion of thanking goodness that one's headache is over, what one thanks goodness for is not that the headache is earlier than another event but *that it is past*. That a future headache is earlier than its cessation doesn't make us dread it less. In terms of discounting, the pain of the headache is discounted as soon as it becomes past.[5] This kind of discounting is intuitively very strong for many objects and experiences.

Future discounting is also typically tensed. While duration – say, a decade – is tenseless, future discounting is usually tensed because the distance is usually measured from the present evaluation point. There is a difference between *wanting something earlier* and *wanting something closer to the present*. In the paradigm temporal discounting experiment, for instance, the subject is asked whether they would prefer a smaller-sooner reward to a larger-later reward. Suppose that they are asked in 2020 about rewards delivered in 2021 and 2022. If they prefer smaller-sooner, say, it is usually not because

[4]Here I'm following the slightly confused tradition of extending the linguistic category of tense to describe mental states like preferences. Syntactically tense is a grammatical device that locates events in the past, present or future. In English, for example, it is often associated with the inflection of verbs, such as *was*, *is* and *will be*. That is analogous to but not the same as the distinction used here, which is probably better captured by the deictic versus sequential time distinction. Arguably, both tensed and tenseless states represent events as standing in a two place temporal ordering relation (earlier, later, simultaneous). That requires two events. "Tenseless" descriptions mention both events overtly. Tensed descriptions mention only one relatum overtly, and allow the other to be filled by a contextual parameter (e.g. the utterance/thought/tokening event). The issue is really whether the event information in the underlying mental representation is overt or covert (and contextually supplied). Viewed in this manner, the tenseless/tensed distinction is just one instance of a far more general phenomenon. See footnote 7 for an example where this point matters.

[5]Here I make no claim about the form of the past discounting function, assuming one can fit the data. It may be that all else being equal I care less about more distant past headaches than less distant past ones, e.g., that I care less about headaches that occurred two years ago than last year. Introspection isn't clear, I think. See Lee et al 2020 and Greene et al 2020.

they have some sort of fixation on the calendar date 2021 and want goods then; rather, it is because 2021 happens to be just one year from now. In the social science literature this is the reason why future discounting is often called delay discounting. How long are you willing to delay the receipt of some goods? The delay is measured with respect to the present evaluation point. The question asked is: how much time *from now* will you wait for goods? That is tensed.

Above I hedge and say "typically" this discounting is tensed because it might be that you do care about goods being delivered on some particular calendar date. Suppose you want flowers delivered on your tenth wedding anniversary. Your care about them declines sharply if delivered afterwards. Discounting like this is perfectly natural. Philosophers tend to disregard such "impure" temporal preferences, ones motivated by the significance of that calendar date and not about time itself. But one could also have a "pure" time preference for a gift arriving on a particular calendar date due to, say, some strange fetish for goods on that date. In the Humean tradition, where preferences are not criticized except when they form inconsistent sets, such a preference is permissible. In both cases the preference is tenseless. The economic framework for discounting can handle tenseless preference, as we'll see; however, the type of discounting one normally studies and cares about in social science is the tensed version. And in philosophy tenseless preferences like these are not viewed as a threat to temporal neutrality, so often they are dismissed.

Since philosophers mostly concentrate on past discounting and ignore distance-dependent features, the tense of temporal discounting is obvious. Because this type of discounting is mostly ignored in social science, however, and because the discussion is dominated by the question of the form of the distance-dependent discounting curve, tense tends to be hidden in this field. Let's now try to find it. Once found, we'll see how it disappears (and maybe returns). By following the trail of tense, we'll see that philosophers and social scientisst don't disagree about temporal discounting as much as initial appearences suggest.

2 The Fate of Tense in Discounted Utility Theory

The standard way of treating future discounting in modern economics arises from Paul Samuelson's 1937 exponential discounted utility theory (EDU). Though the model has been controversial for almost its entire life, it is still commonly used in cost benefit analyses and can fairly be considered the mainstream account of temporal discounting. Samuelson explicitly considered and dismissed a normative interpretation of EDU. He did not think one *should* discount according to EDU; he mainly felt that simplicity and usefulness recommended it. Another founder of EDU, Ramsey 1928, also rejected a normative understanding of its permitted temporal discounting.

It wasn't until the seminal work of Strotz 1955 that EDU gained its normative interpretation. Strotz shows that if one discounts according to any non-EDU-endorsed schedule, one can suffer preference reversals that make one in principle exploitable. Because being in principle exploitable is the cardinal sin of modern economics and rational choice theory, Strotz's result placed EDU on a normative pedestal, a place where it has remained

(even if Strotz himself advanced a non-EDU account).

Work in psychology and behavioral economics challenge this framework by showing that it isn't descriptively accurate. People don't actually discount according to EDU. Given the normative standing accorded EDU, people are thus interpreted as yielding to irrational preferences, preferences that predictably lead to poor individual and social outcomes. All of this is controversial, but it is the conventional understanding.[6]

How does EDU treat time? Discussions are sometimes fairly coarse, failing to make important distinctions. Fortunately, Strotz himself is admirably clear, so we'll begin with this canonical source.

2.1 Strotz, Time and Consistency

Discounted utility theory is an extension of expected utility theory that seeks to incorporate time preferences. EDU is a particular model of discounted utility theory, one that picks a special form for modelling time preference. In what follows, let me describe the general landscape of discounted utility theory and then show how Strotz lands on EDU as the normatively correct model.

When navigating through life, you face many choices with outcomes distributed throughout time: (e.g.) saving for retirement, suffering through a dentist visit, going to surf. Following Strotz, think of these different paths as curves in a consumption-time diagram. Each curve $C(t)$ specifies a particular set of choices through some period $0 \leq t \leq T$. We assume that your preferences are sharp enough to let you transitively order all the different alternatives you face, $C_1(t)$, $C_2(t)$, $C_3(t)$... and that this ordering can be represented by a utility functional Φ_τ. The subscript τ represents the evaluation time. We can think of it as the present. Φ_τ specifies your ranking at the present time $t = \tau$ of consumption paths through time .

Discounted utility theory is a simple extension of expected utility theory. To focus on temporal discounting, it's convenient to ignore the "expected" part of expected utility theory and pretend that everything is certain. We also don't want to have to trace through time all the knock-on effects of choosing goods, experiences, and so, such as the pleasant memories a vacation make provide. We'll therefore focus on instantaneous utility. The utility function $u[C(t), t]$ thus assigns a value $u(t)$ to $C(t)$ for every time t.

To represent discounting Strotz proposes the function $\lambda(t - \tau)$. The present moment is given by τ and time t is some date in the past or future of τ. The discount function is thus a function of the distance between that date, say, New Year's Day 2030, and the present. Since we don't discount the present, we assume that $\lambda(0) = 1$.

We can now state the goal of discounted utility theory, which is to maximize the utility functional

[6] There are many authors who claim for a variety of different reasons that non-EDU discounting can be rational; see in philosophy (e.g.,) Callender 2021a, Pettigrew 2020, and in social science (e.g.) Burness 1976, Drouhin 2009, McGuire and Kable 2013.

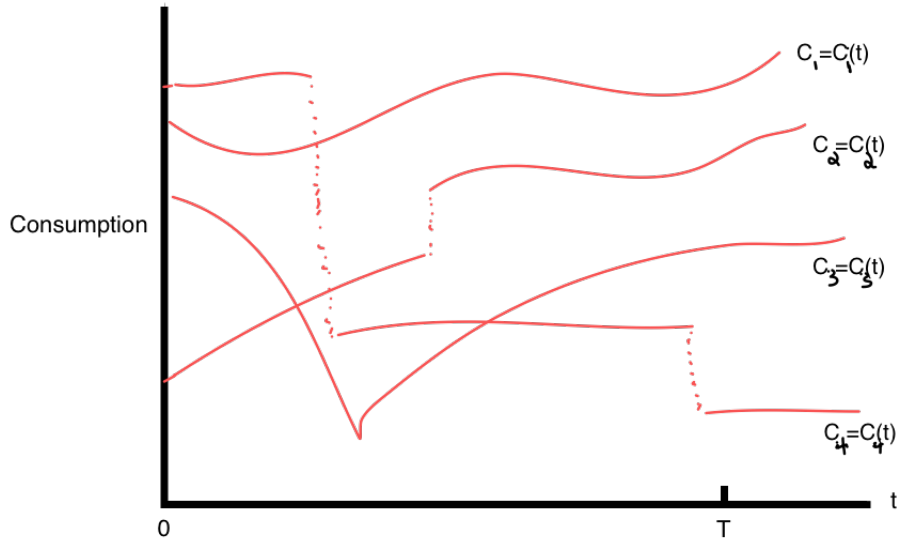

Figure 1: Consumption paths through time

$$\Phi_\tau = \int_0^T \lambda(t-\tau)u[C(t),t]dt \tag{1}$$

where the integrand factors into two functions, a discount function of $(t-\tau)$ and a utility function of t.

Strotz carefully notes that (1) may depend on two types of discounting, discounting from the present moment or from the calendar date of the future consumption. For the latter, he gives the example of wanting champagne on his birthday; receiving it afterwards is of much less value. This kind of (what philosophers will call) "impure" temporal preference gets encoded in the utility function and not the discount function, for the weight $u(t)$ assigned to the pleasure of champagne is a function of calendar time t. The value of the champagne will then be high when t=birthday and decline thereafter. Equation (1) captures both types of time preference.

We wish to find the optimal plan forward from the present $\tau = 0$. Assume that we have a finite stock of some resources to be consumed. We want to maximize (1) subject to constraints imposed by this stock, beginning at $\tau = 0$, with endpoints given by $t = 0$ and $t = T$. Plugging in these constraints and sparing the math, the solution to this differential equation is

$$\lambda(t) \cdot \frac{\partial u[C(t),t]}{\partial C} = constant$$

where the constant is dependent upon the stock. In words, what this means is that *spreading the discounted marginal utility of the stock equally across time* maximizes utility. If our stock is water, we would have to take account of the diminishing utility of extra increments of water – the first glass quenches thirst but later glasses aren't as valuable – and our time preference that makes a glass tomorrow less valuable than one tomorrow. An optimal planner spreads this product evenly across time.

This result is restricted to a plan as seen from the present. What about consistency across time? In this case we wish to find the optimal path through time at $\tau = 0$, reevaluating at some new present τ_1, and then another τ_2, and another, for all τ. Strotz asks, "Under what circumstances will an individual who continuously re-evaluates his planned course of consumption confirm his earlier choices and follow out the consumption plan originally selected?" (171).

Again sparing the math, the answer is that consistency is achieved when one discounts according to the exponential discount function

$$\lambda(t) = k^t \tag{2}$$

where k is a constant and $0 \leq t \leq T$. (2) makes one's discounting constant and proportionate to amounts of time. See Figure 2. The function is none other than the simple discount function chosen by Samuelson. The model of discounted utility theory that incorporates the exponential discount rate (2) is called *exponential discounted utility theory* (EDU). Strotz proves that (2) is the unique function that will lead to time consistent choices. Any other may lead to preference reversals and possible exploitation.

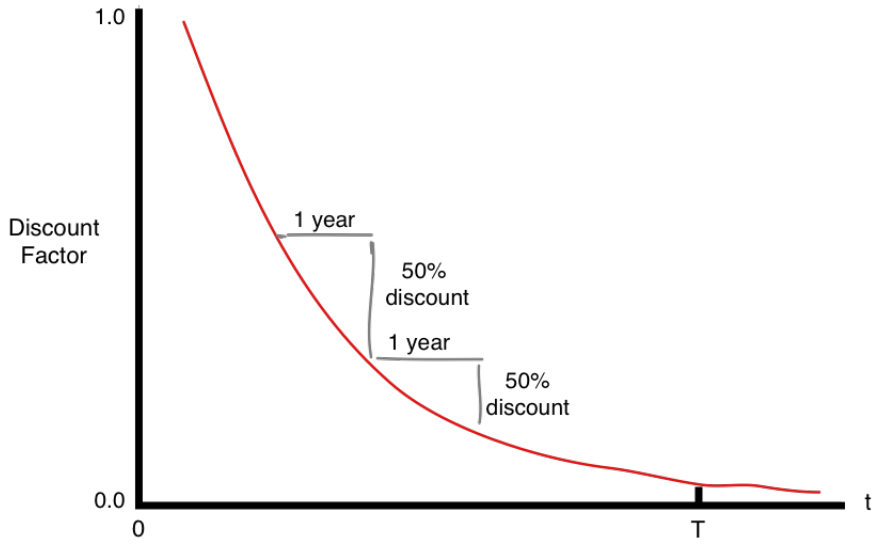

Figure 2: Exponential discounting. The amount discounted is proportional to the amount of calendar time.

Plugging (2) into (1), observe that the present time τ has vanished.

2.2 Tense and Strotz's Result

The parameter τ represents tensed time in discounted utility theory. It stands for the time of evaluation, the present moment for the decision-maker. The parameter t by contrast runs over a timeline that represents earlier and later than relations amongst events. Think of the timeline ordered by t as a map. Then τ is the red dot indicating "you are (temporally) here." In the language of philosophy of time, τ is an A-determination and a particular instance of t is a B-determination.

Having recognized this, what immediately jumps out at us is the fact that equation (2), the backbone of normatively charged EDU, lacks any reference to τ, and hence, tensed time. When comparing two outcomes to be delivered at different times, the only things that matter for an EDU discounter are (i) the tenseless distance between the two outcomes and (ii) the different valuations of the outcomes. That's it. Both (i) and (ii) are entirely tenseless. Tense doesn't matter at all when you are making a decision according to EDU.

This result is at once surprising and not surprising at all. On the one hand, it's surprising because discounting is a thoroughly tensed affair. Except when discounting for calendar date, what matters to the decision-maker is the delay from now. It's alarming that the "now" drops out. On the other hand, it's not surprising that it does. That was more or less the goal of the enterprise. Finding a rate that one will reaffirm, again and again, is very close if not tantamount to finding a rate that doesn't depend on when is now. Strotz's question demands that you stand by your evaluations in every context, so it's unsurprising that an evaluation that changes its truth-value depending on a contextually supplied parameter (e.g., τ) is not permitted.

To be clear, EDU doesn't *do away* with tense. We can't accuse Strotz of taking tensed discounting and then proving a result about something different–tenseless discounting. The discounting he works with is explicitly tensed. It's just that the tense drops out when temporal consistency is demanded. The discount rate can make reference to a now, but it has to be independent of when now happens to be. Tense isn't ignored. It is forbidden to matter in the representation of the answer due to the nature of the question.

I think it's illuminating to view this central result in economics and rational choice theory via tense. The result is saying tensed discounting is rational when its tense doesn't matter. This interpretation will help us connect EDU to debates about the philosophical position known as temporal neutrality.

Two other features of this result as it relates to tensed time are interesting.

First, the utility function $u[C(t), t]$ is solely a function of t and not τ. This is highly significant. In the philosophical literature, many famous examples turn on tense mattering to how much we value an outcome. In Parfit's famous examples, we're willing to trade a small amount of future pain for more past pain. The intuition is that past pain is "over and done" and therefore not something we care as much about as future pain. With u not being a function of τ, we cannot represent this kind of change in value/care as a change in utility. If you recall, Strotz gave the example of wanting champagne on his birthday,

which is some calendar time t. The utility conferred by the champagne can depend on t, so birthday-champagne can be valued more than day-after-birthday champagne. A similar dependence cannot happen for changing tensed perspective via τ, however, so champagne-on-birthday in the past is valued the same as champagne-on-birthday in the future.[7]

Second, the integrand (1) factors into two functions, a function of the temporal distance from now and a function of only calendar time. This means that there can be no interaction effects between tensed discounting and calendar discounting. Yet arguably there are. Maybe in 1980 when I was young I cared less about the past than I do now in 2020; or perhaps in 2020 I will discount the future less sharply than in 2030 because it matters more to me then.

3 EDU and the Past

The reader may wonder whether EDU can treat past discounting. The intuitive pull of past discounting can be strong. In Parfit's famous case, one wakes up in a hospital and is told that either one has undergone a very painful surgery while conscious yesterday or will undergo such a surgery tomorrow. You cannot remember yesterday due to an amnesia pill you hav etaken. Yesterday's surgery, if it happened, lasted for X hours; tomorrow's, if it happens, will last for Y hours. When asked, some people prefer to find out that they already had the surgery, even if $X >> Y$ (see Lee et al 2020). Can EDU model this kind of preference?

In Strotz's model we are looking at some period $0 \leq t \leq T$ and then assuming that the first decision point is at $\tau = 0$. So the whole time period covered is the future. This treatment is very typical in economics. However, it is certainly possible to treat the past in this model, and indeed, Strotz does. We can allow t to range over all of time and pick τ to have any value of t we desire. Suppose we discretize time into years and let the "first" moment of time be $t = 0$ and our first decision be 13 billion years later at $\tau = t_{2021} = 2021$. Then all $t < \tau$ is "the past."

We now have a past but also a big problem. An exponential discounter will, for times $t < \tau$, value the past more than the present, and value the far past more than the near past (Hedden 2015 also makes this point). That isn't even close to descriptively accurate, as that is more or less the opposite of what we do. In fact, it's not past discounting at all, but past inflation.

Perhaps one might impose a rule demanding exponential discounting *away* from time $t = \tau$ in *both* temporal directions. Our concern would peak at $t = \tau$ and then fall away exponentially toward the past ($t < \tau$) and future ($t > \tau$). Ignoring fit with actual data, that at least would allow for past discounting in the model. Here we would be taking advantage of the fact that Strotz's model allows for future inflation as much as future

[7] This kind of phenomenon is true of non-temporal essentially indexical utilities too. I care more about whether *I* get champagne than whether Craig Callender gets champagne. I care much more whether the champagne is delivered *here* than whether it is delivered to the particular address in San Diego County where I happen to be. And so on.

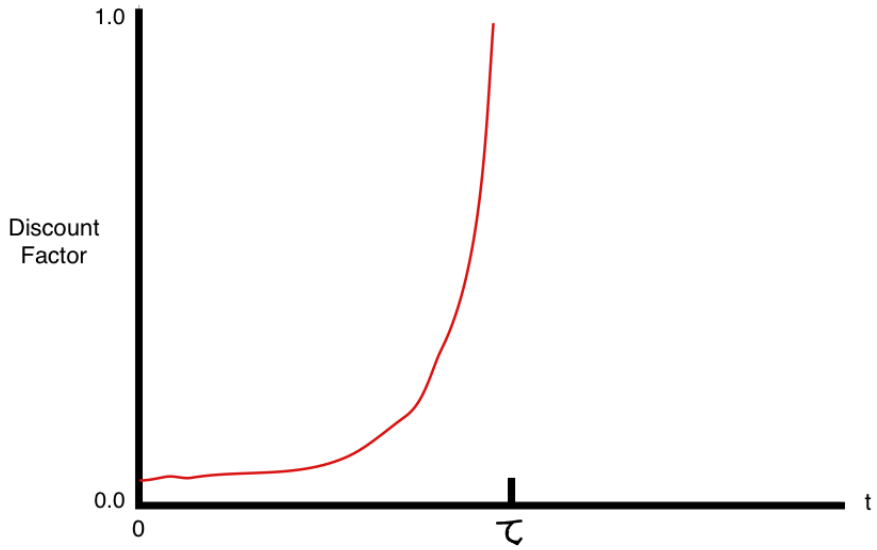

Figure 3: Past discounting

discounting: a curve that inflates the future, when shifted into the past, will discount the past. The problem with this modified theory, however, is that it will face the same problems non-exponential discounting did. This sharp "peak" about the present moment would be non-exponential; and as the present moment updates, the discounting will not be time consistent. The upshot is that although past discounting can be modeled in discounted utility theory, there is no way to do so in a time consistent manner because it is tensed.

Economics ignores past discounting. The common reason mentioned for this practice is that the past is uncontrollable so we don't need to consider decisions about it. That is undoubtably true for most decisions considered in economics, but it is not as cost-free a restriction as is sometimes thought. In philosophical thought experiments, for instance, Doughtery 2011 shows that in concert with your other preferences – specifically, risk aversion – past discounting can matter to the future. And more generally, past discounting can matter in all sorts of practical matters. For instance, consider price negotiation. Taking advantage of someone's past discounting, I may be able to negotiate a lower price for a service already rendered than one to be rendered. Still, since these are not the sorts of problems for which economists use discounted utility theory, few will lose sleep if they must ignore past discounting.

4 Two Types of Temporal Neutrality and EDU

The philosophical position of temporal neutrality can be traced back to ancient times. In the modern period, one finds a clear expression in Adam Smith 1790. Smith holds that the prudent person heeds the approbation of an "impartial spectator":

> 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 the same: he sees them nearly at the same distance, and is affected by them very nearly in the same manner.(VI.i.11)

Temporal neutrality is explicitly tied to rationality by Henry Sidgwick 1907, who writes that

> Indeed this equal and impartial concern for all parts of one's conscious life is perhaps the most prominent element in the common notion of the rational–as opposed to the merely impulsive–pursuit of pleasure. (111)

and this position is endorsed by John Rawls, e.g., "The mere difference of location in time, of something's being earlier or later, is not in itself a rational ground for having more or less regard for it" (Rawls 1971, 293-4). David Brink provides a recent succinct statement of the position:

> temporal neutrality should be understood to claim that the temporal location of goods and harms within a life has no normative significance except

> insofar as it contributes to the value of that life. We might say that on this view temporal location has no independent significance or no significance per se. (358)

Advocates of temporal neutrality are not always clear about what kind of time series – a tensed A-series or tenseless B-series – they mean when they refer to "temporal location." Location in what series? In the above quotes, Smith picks out A-properties (present, future) whereas Rawls mentions B-properties (earlier, later). Rawls in the very next sentence (not quoted) then switches to A-series language. This imprecision occurs naturally in ordinary language depending on what is assumed in context. None of the thinkers mentioned are concerned with temporal linguistics, so this kind of shifting back and forth is perfectly natural. However, making the distinction allows us to distinguish two types of temporal neutrality, one that denies that tensed location has independent significance and one that denies that tenseless location has independent significance.

The tensed reading, I think, is the most natural understanding of temporal neutrality. The motivation for the thesis in Smith, Sidgwick and most others is to advise people not to give in to impulsive acts that satisfy the momentary present self. Instead the temporal neutralist urges us to consider one's overall good across one's whole life. Much of this literature is best interpreted as against what psychologists often call "present bias". Temporal neutralism's recommendation of pursuing now-for-later sacrifices is thought to be needed to counteract this tendency to inflate present value. It is not best interpreted as counseling, say, 2025-for-later sacrifice, except when the now occurs in 2025. Strotz speaks of the "intertemporal tussle" in which we all engage, the tension between what the momentary self wants and what is good for the whole person extended throughout time. Temporal neutralism is naturally understood as taking the side of the whole temporally extended person. In so doing it rejects the importance of satisfying any particular time slice of you except insofar as such satisfaction can be justified for the person as a whole.

Further evidence for thinking the tensed reading is standard comes from examining a sophisticated neutralism, such as Brink's, which allows that one may desire one's life to have certain temporal patterns. For instance, maybe you prefer a rags-to-riches life to a riches-to-rags life (Velleman 1991). If so, then you prefer to distribute your resources toward the later moments of your life than the earlier. You're willing to sacrifice earlier-for-later more than later-for-earlier. This pattern is entirely compatible with neutralism. The reason why they're compatible, I would say, is that these second-order temporal preferences are all tenseless temporal preferences. They are condoned by Smith's impartial spectator looking over your entire life. What is anti-neutralist is a preference that depends upon where one is presently located in one's lifetime. Where the "red dot" is on your worldline should not matter if you're a neutralist. In other words, tense should not matter.

A tenseless reading is also possible. In this case the temporal location that has no significance is a location in the B-series. This reading is fine but it does make temporal neutrality trickier to defend than the tensed variety. The reason is that it's very hard to disentangle B-location and non-temporal significance. B-location almost always matters: one location is your birthday, in another you're young, in another you're old, in another

you're in a different country, and so on. That it's okay to discount the value of receiving champagne after your birthday has passed is agreed to be acceptable by everyone. To find the controversial discounting, one needs to abstract away from all of these "impurities" and isolate cases of B-theoretic discounting that are pure. What one is left with is a person who has a kind of fetish for a particular calendar date, like a tenseless version of Parfit's example of a man who, all else being equal, prefers his pain on Tuesdays. (Parfit's example is about future Tuesdays, not later Tuesdays.) One can certainly object to this kind of preference. But this question really has more to do with how "Humean" one is about the permissibility of preferences and less about time per se. If temporal neutrality is about this kind of discounting being irrational then it's hard to see what all the fuss is about – there aren't many of these "ideally coherent eccentrics" (Street 2009) walking around.[8]

EDU can model both types of discounting. We've already seen that Strotz encodes tensed discounting in the temporal discounting function and represents tenseless discounting by allowing the utility to be a function of calendar time t. Suppose we accept that, however one defines temporal neutrality, its core commitment is to tensed time not having significance. Viewed this way, we can see that EDU builds in a kind of temporal neutrality. It sanctions those temporal preferences that cannot be exploited and that maximize utility along a lifetime. And in particular, as we saw, tense drops out of optimal discounting.

Perhaps surprisingly, this allows some future discounting. Suppose you're offered a small reward now or a larger reward next week. Should you take the smaller sooner reward? EDU doesn't condemn you if you do. It doesn't judge preferences. But it does say that if you do accept the more immediate reward, your position on your timeline shouldn't be relevant. So if you take the smaller sooner reward now, you should also take the smaller sooner award if the choice is between rewards spaced one year from now and a year and a week from now. The discounting is only a function of two differences, (big reward - small reward) and (late date - soon date). If you discount, it's okay so long as you always do it the same way. "When" you are therefore doesn't matter. This is a kind of discounting that an ideal spectator could endorse.

We now see that EDU is more or less committed to temporal neutralism being the normatively correct story. The Appendix goes through the logic linking temporal neutralism and EDU (and see Callender 2021b for more detail and discussion, including a derivation of EDU from temporal neutrality). But the main point is that what at first looked like a tension between many philosophers and economists is in fact a consilience.[9]

[8] Greene 2021 argues that preferences partly based on pure time preferences are objectionable. In this way he hopes to avoid disentangling pure from impure preferences and make his recommendations widely applicable. See Callender 2021b for some thoughts on this maneuver.

[9] I've spelled out temporal neutralism in two ways that I consider natural, neutral about position in the A-series and neutral about position in the B-series. That doesn't mean that there aren't other ways of doing so that are strong enough to eliminate any type of future discounting whatsoever. For instance, one might insist on being neutral about the relative delays from the evaluation point, so that if one preferred reward x to y when spaced apart by temporal distance Δt, then one prefers x to y no matter what $\Delta t'$ is. That would eliminate discounting altogether and regain a disagreement between some economists and philosophers. Still, I've shown that on at least once natural way of

5 Hyperbolic Discounting and Tense

Strotz states that people are not actually exponential discounters. With great prescience, he develops models of how we do discount that are now called *hyperbolic discounting*. Many decades after Strotz, with interest from psychology and the rise of behavioral economics, empirical study after empirical study demonstrated conclusively that we are not exponential discounters. We are said to be "hyperbolic" discounters instead because some of the models suggested to fit the data include a discount function with a hyperbolic or quasi-hyperbolic form. Scores of different "hyperbolic" discounting models have been proposed.[10] None fit all the data, but each typically has some kind of rationale (e.g. the uncertainty of hazards) accompanying it that attempts to explain why decision-makers would employ it.

A kind of familiar narrative emerged. Dual process theory developed in the heuristics and biases tradition of work on judgements. That tradition, begun by Tversky and Kahneman (see e.g., Tversky and Kahneman 1981), showed that we often violate expected utility theory in our reasoning; for example, we commonly commit probabilistic fallacies. Dual process theory arose as the explanation for these "anomalous" departures from the normative standard of reasoning. As the name suggests, it proposes two types of mental processing, system 1 and system 2. System 2 is our cool, rational system. It is typically slow, reflective and conscious. System 1 is our "lower" more intuitive system. It is typically connected to the emotions, fast, heuristic, and unconscious. The exponential versus hyperbolic debate was fit to this narrative (see, e.g., Loewenstein and O'Donoghue 2005). Exponential discounting is seen as the work of rational system 2 processing, whereas the present bias of emotional system 1 is understood as bending the exponential curve toward the hyperbolic in greedy service of immediate needs.

Challenges to this picture have arisen from many quarters. As mentioned, many disagree that EDU is the right normative standard for discounting, but others have also questioned whether dual process theory is a suitable expanatory model (see, e.g., Samuel 2009, Pennycook 2017).

It's eye-opening to view this debate through the lens of tense.

As mentioned, there are scores of different hyperbolic functions proposed. Many have remarked that "hyperbolic" discounting is better characterized as simply "non-exponential" discounting. But the differences amongst these functions won't matter much for us. Consider the functional form of the simplest and most canonical of forms

$$\lambda(t-\tau) = (1+\alpha(t-\tau))^{-\beta} \tag{3}$$

where α and β are constants. We notice immediately that τ is back![11]

understanding temporal neutrality, the dispute goes away.

[10] See Thaler 1981, Loewenstein and Prelec 1992, Laibson 1997, Loewenstein and Donoghue 2005 and Sozou 1998 for a start.

[11] Not seeing this function always written this way, the reader may suspect that it is back only because I put it back. It's true that some presentations of hyperbolic discounting present a function like that in (3) lacking the τ. However, when actually used, careful discussions of hyperbolic discounting always distinguish between calendar (or date) time and delay time, and it is not at all controversial that

Tense is crucial to hyperbolic discounting. Let's see how this works with a simple example. Suppose that you don't like cleaning the gutters of your house. You're wondering when to do it, where in the example time is discrete and the units are days. You do not discount today, $\lambda(0) = 1$, but you do today care less about its cost tomorrow, i.e., $\lambda(1) < 1$. If you were an exponential discounter, then we know that the tensed day doesn't matter to you. Hence your temporal preferences satisfy

$$\frac{\lambda(0)}{\lambda(1)} = \frac{\lambda(365)}{\lambda(366)}$$

because a day's wait is a day's wait, no matter when it happens. If you are a hyperbolic discounter, the tensed day does matter to you; in fact, it can matter quite a lot. Let $\alpha = \beta = 1$. You still do not discount today, so $\lambda(0) = 1$. But you discount cleaning the gutters tomorrow by $\lambda(1) = (1 + 1)^{-1} = 1/2$. You dislike cleaning the gutters today twice as much as doing it tomorrow because $\lambda(0)/\lambda(1) = 2$. Because of the presence of τ in (3), the ratio between cleaning the gutters a year from now versus a year and a day from now is not the same

$$\frac{\lambda(0)}{\lambda(1)} \neq \frac{\lambda(365)}{\lambda(366)}$$

and in fact you don't much care between these two distant days because $\lambda(365)/\lambda(366) = 1.003$. Your hatred of gutter cleaning flattens out from today.

When put in terms of something valued (say, a cash reward) versus disvalued (say, cleaning the gutters), this "flattening out" is characterized as diminishing impatience. Diminishing impatience is often regarded as the central feature of hyperbolic discounting. And indeed, it is important, for one wants to fit the empirical data that displays such diminishing impatience. What is absolutely crucial to getting the intended "hyperbolic" behavior, however, is the inclusion of the present, τ. The tensed present is the source of the famous violations of temporal consistency. When we fast forward a year, you will view cleaning the gutters tomorrow (what used to be Day 366) as twice as good as cleaning them today (what used to be Day 365), not merely $1.003x$ better. A clever gutter cleaning service might in principle be able to exploit the changing values of $\lambda(365)/\lambda(366)$ you suffer as these dates draw near. Whether this temporal inconsistency is really a sign of irrationality is a matter of great debate, as mentioned.[12] Rational or not, since people do display this pattern of behavior, descriptively adequate models need to incorporate tense.

As Rasmussen 2008 points out, one can be a temporally consistent yet non-exponential discounter so long as (in my terminology) the discounting is tenseless. Suppose that you

hyperbolic time essentially employs both. As Rasmusen 2008 notes, where his "relativistic" time is my tensed time, "Hyperbolic discounting makes not one but two changes from the standard model. First, it makes the per-period discount rate change over time. Second, it bases discounting on relativistic time rather than absolute time. It is this second assumption which is the key one."

[12] For one, the claim that it is irrational assumes that nothing happened in that intervening time that allows you to rationally change your mind (it assumes that a condition know as Invariance (see Appendix) holds and that it is normatively warranted). See Callender 2021a, Halevy 2015, Janssens, Kramer, and Swart 2017, and Pettigrew 2020.

discount at a rate of 5% during the decade [2029, 2039] and then switch to 10% during [2040, 2050]. Perhaps you wish to discount more when you are older. This discounting is non-exponential because non-constant. So long as it is understood as about tenseless calendar dates, however, it causes no problem or inconsistency. Cleaning the gutters in 2038 gets discounted at 5% and cleaning them in 2045 gets discounted at 10%. But if I know that now, I can calculate those values, plug them into my utility function, and generate no inconsistency so long as I stick to my plan and keep cleaning the gutters in those two years discounted as I did originally. Rasmussen puts his finger precisely on what matters:

> The key is that exponential discounting treats the parameter as "Rasmusen's rate of time preference for when he is 70 in the year 2058" whereas non-exponential discounting treats it as "Rasmusen's rate of time preference for 21 years from the present."

What matters is tense. Tense is what allows the discounting of a future event to change as we approach it.

This point explains what might otherwise be confusing: the pheneomenon of time-consistent hyperbolic discounting. Drouhin 2015 and Burness 1976 emphasize that contrary to conventional wisdom hyperbolic discounting can be consistent. Their trick is to make calendar time and evaluation time multiplicatively separable in the discount factor. Without going into the details, this leads to a discount rate that declines in calendar time t and is *independent of the evaluation time*. The same goes for Weitzman 2001's famous use of a declining social discount rate. Both types of model employ non-constant discounting but avoid inconsistency because their hyperbolic functions aren't a function of evaluation time.[13] Put in my terms, they don't lead to preference reversal because they are not a function of tense. Constant discounting leads to time consistent exponential discounting, but that doesn't mean non-constant discounting leads to time inconsistency: crucially one needs that function to also be a function of tensed perspective.[14]

The great debate between exponential and hyperbolic discounting is thus a debate between tenseless and tensed discounting. As philosophers have noticed with intuition pumps and as behavioral economists and psychologists have empirically demonstrated, the folk discount in a tensed way. Temporal neutrality and the desire to maximize utility over a lifetime and not at a time, however, lead some philosophers and economists to hold that tenseless discounting is the only one that can be rational. The debate within the hyperbolic discounting camp is about the form of the function; but the larger debate is really over whether this function is tensed or not. In this debate on the rationality of

[13] The discount rate in Drouhin 2015 and Burness 1976 is $\rho(t) = {}^{\alpha\beta}/_{(1+\alpha t)}$ where α and β are constants and t is calendar time. Weitzman's discount rate aggregates individual rates that disagree. His function can be written as $\rho(t) = {}^{\mu}/_{1+\frac{\sigma^2}{\mu}\cdot t}$ where μ is the mean and σ the standard deviation of individual discount rates. What's important for us is that although hyperbolic and declining, the evaluation time τ drops out in both cases, just as it does in Strotz's result.

[14] In the terms used in the Appendix, these models maintain consistency by violating stationarity and invariance. The violation of stationarity is what makes them hyperbolic, and the violation of invariance means that they "care" about calendar time.

future discounting we see an echo of the debate in philosophy over the rationality of past discounting. Both are debates about the rationality of discounting for tense.

6 Conclusion

This paper hasn't resolved any of the big debates regarding tense and rationality. I hope that it's made progress, nonetheless, by bringing the social science and philosophy on temporal preferences closer to one another. On its face, it looks like work in social science and philosophy are in sharp disagreement over the rationality of the two different types of temporal discounting. But when we peer at the debates through the lens of tense, we find more in common than first meets the eye. EDU is a version of temporal neutralism. In fact, it is essentially implied by the conjunction of expected utility theory and temporal neutrality (see Appendix; Callender 2021b). Temporal neutralists often take themselves to be against temporal discounting; but we've seen that EDU allows only a kind of tenseless discounting, a type of discounting that arguably is consistent with the spirit of neutralism. Regarding past discounting, philosophers divide on whether it is rational. Temporal neutralists think it is not whereas others side with the strength of the intuition. EDU, we've seen, doesn't really treat past discounting, but if one tries to incorporate it then it won't be judged as rational. So EDU sides with temporal neutralists on past discounting, as it more or less does on future discounting. Just as philosophy divides on past discounting, we saw that social science divides on the rationality of EDU. The great debate between exponential and hyperbolic discounting, we saw, is very much a dispute over the rationality of discounting for tense.

By drawing past and future discounting together under the umbrella of tense, I am not suggesting that they are the same sorts of thing. They are both tensed; through this lens the debates on rationality in social science and philosophy line up closer than ordinarily conceived. In this important way they are alike. But I make no claim about whether they stem from the same source, whether they have the same kind of explanation, and so on. Although I suspect they are deeply linked via the situated momentary self (Ismael 2017, Callender 2017), I believe that they are otherwise quite different in many respects. Do they treat hedonic and non-hedonic experiences alike? Positive and negative valanced events? Experienced and unexperienced events? Does each asymmetry extend to the third person? Are they associated with emotion in the same way? As many of the chapters in this volume demonstrate, psychologists and experimental philosophers are just beginning to probe these two asymmetries. What they've found so far is that our discounting behavior is extraordinarily complex. Given the vast differences between the future, which is viewed as inherently risky, for instance, and the past, which is not viewed as risky, it's hard to believe that this web of preferences and behavior will turn out to be the same for the two types of preferences.

That said, now that we can see clearly just how important tense is and how both social science and philosophy are debating the rationality of discounting for tense, this paper does suggest that future research on the origin of our tensed model of the world may be important. Parfit famously tried to justify tensed (past) discounting via a tensed

metaphysics. He had a very hard time doing so.[15] Perhaps if we better understood the role of tense , however, we could make more progress.[16]

References

[1] Ahmed, A. 2018. Rationality and Future Discounting. *Topoi*, 1-12.

[2] Brink, D. 2011. Prospects for Temporal Neutrality. In Callender, C. (ed) *The Oxford Handbook of Philosophy of Time*. Oxford: Oxford University Press, 353-381.

[3] Burness, H. S. 1976. A Note on Consistent Naive Intertemporal Decision Making and an Application to the Case of Uncertain Lifetime. *Review of Economic Studies* 43(3), 547-549.

[4] Bykvist, K. 2015. Value and Time. In Hirose, I., and Olson, J. (eds), *The Oxford Handbook of Value Theory*. Oxford University Press, 117-135.

[5] Callender, C. 2017. *What Makes Time Special?* Oxford: Oxford University Press.

[6] Callender, C. 2021a. The Normative Standard in Future Discounting. *Australian Philosophical Reviews*, forthcoming.

[7] Callender, C. 2021b. A Derivation of Exponential Temporal Discounting from Temporal Neutrality, forthcoming.

[8] Caruso, E., Gilbert, D. and Wilson, T. 2008. A Wrinkle in Time: Asymmetric Valuation of Past and Future Events. *Psychological Science* 19, 796–801.

[9] Dougherty, T. 2011. On Whether to Prefer Pain to Pass. *Ethics* 121, 521-537.

[10] Dorsey, D. 2017. Future-Bias: A (Qualified) Defense. *Pacific Philosophical Quarterly* 98, 351– 373.

[11] Drouhin, N. 2009. Hyperbolic Discounting May Be Time Consistent. *Economics Bulletin* 29, 2552–2558.

[12] Fernandes, A. 2019. Does the Temporal Asymmetry of Value Support a Tensed Metaphysics?. Synthese, forthcoming. https://doi.org/10.1007/s11229-019-02269-8

[13] Fishburn, P. and A. Rubinstein, 1982. Time Preference. *International Economic Review* 23, 677-694.

[14] Greene, P., Latham, A., Miller, K. and Norton, J. 2020. Hedonic and Non-Hedonic Bias toward the Future. Australasian Journal of Philosophy. 1-16.

[15] See Hare 2008 for some problems.

[16] Many thanks to David Brink, Jennifer Carr, Jonathan Cohen and Alison Fernandes for useful comments and discussion.

[15] Greene, P., Latham, A., Miller, K. and Norton, J. 2021. Why are People so Darn Past Biased? This volume.

[16] Greene, P. 2021. 'Pure' Time Preferences Are Irrelevant to the Debate over Time Bias: A Plea for Zero Time Discounting as the Normative Standard. Response to Callender. Australasian Journal of Philosophy, forthcoming.

[17] Halevy, Y. 2015. Time Consistency: Stationarity and Time Invariance. *Econometrica* 83, 335–352.

[18] Hare, C. 2008. A Puzzle about Other-directed Time-bias", *Australasian Journal of Philosophy* 86, 269–277.

[19] Hedden, B. 2015. *Reasons without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.

[20] Ismael, J. 2017. Passage, Flow, and the Logic of Temporal Perspectives. In Christophe Bouton and Philippe Huneman (eds.), *Time of Nature and the Nature of Time*. Springer Verlag, 23-38.

[21] Janssens, W., Kramer, B., and L. Swart. 2017. Be Patient When Measuring Hyperbolic Discounting: Stationarity, Time Consistency and Time Invariance in a Field Experiment. *Journal of Development Economics* 126, 77-90.

[22] Koopmans, T. C. 1960. Stationary Ordinal Utility and Impatience. *Econometrica* 28, 287- 309.\

[23] Laibson, D. 1997 Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112, 443-477.

[24] Lee, R. Hoerl, C. Burns, P., Fernandes, A., O'Connor, P. and T. McCormack. 2020. Pain in the Past and Pleasure in the Future: The Development of Past-future Preferences for Hedonic goods. *Cognitive Science* 44(9), forthcoming.

[25] Lee, R. and McCormack, T. 2021. A Developmental Perspective on Temporal Asymmetries, this volume

[26] Loewe, G. 2006. The Development of a Theory of Rational Intertemporal Choice. *Revista de Sociologia* 80, 195–221.

[27] Loewenstein, G. and Prelec, D. 1992. Anomalies in Intertemporal Choice: Evidence and Interpretation, *Quarterly Journal of Economics* 107, 573-597.

[28] Loewenstein, G., and O'Donoghue, T. 2005. Animal Spirits: Affective and Deliberative Processes in Economic Behavior. Working Paper 04-14, Center for Analytic Economics, Cornell University.

[29] Maclaurin, J., and Dyke, H. 2002. 'Thank Goodness That's Over': The Evolutionary Story. *Ratio* 15, 276–292.

[30] McGuire, J. and Kable, J. 2013. Rational Temporal Predictions Can Underlie Apparent Failures to Delay Gratification. *Psychological Review* 120, 395–410.

[31] McTaggart, J.M.E. 1908. The Unreality of Time. *Mind* 17, 457-74.

[32] Parfit, D. 1984. *Reasons and Persons* (Vol. II). Oxford: Clarendon Press.

[33] Peart, S. 2000. Irrationality and Intertemporal Choice in Early Neoclassical Thought, *Canadian Journal of Economics* 33, 175-189.

[34] Pennycook, G. 2017. A Perspective on the Theoretical Foundation of Dual Process Models. In De Neys, W. (Ed.), *Dual Process Theory 2.0*, 13–35.

[35] Pettigrew, R. 2020. *Choosing for Changing Selves*. Oxford University Press.

[36] Prior, A.N. 1959. Thank Goodness That's Over. *Philosophy* 34, 12 - 17.

[37] Ramsey, F.P. 1928. A Mathematical Theory of Saving. *Economic Journal* 38, 543-549.

[38] Rasmusen, E. 2008. Some Common Confusions about Hyperbolic Discounting. Working Papers 2008-11, Indiana University, Kelley School of Business, Department of Business Economics and Public Policy. Available at SSRN: https://ssrn.com/abstract=1091392.

[39] Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.

[40] Samuels, R. 2009. The Magical Number Two, Plus or Minus: Dual Process Theory as a Theory of Cognitive Kinds. In Keith Frankish & Jonathan St B. T. Evans (eds.), *In Two Minds: Dual Processes and Beyond*. Oxford University Press, 129-146.

[41] Samuelson, P. 1937. A Note on Measurement of Utility. *Review of Economic Studies* 4, 155-161.

[42] Sidgwick, H. 1874. *The Methods of Ethics*. London: MacMillan.

[43] Smith, A. 1790. *The Theory of Moral Sentiments*. Oxford: Oxford University Press.

[44] Sozou, P.D., 1998. On Hyperbolic Discounting and Uncertain Hazard Rates. *Proceedings of the Royal Society B: Biological Sciences* 265, 2015–2020.

[45] Story, G. W., Vlaev, I., Seymour, B., Darzi, A., & Dolan, R. J. 2014. Does Temporal Discounting Explain Unhealthy Behavior? A Systematic Review and Reinforcement Learning Perspective. *Frontiers in Behavioral Neuroscience* 8, 76.

[46] Steele, K. 2021. Why Time Discounting Should be Exponential: A Reply to Callender. *Australian Philosophical Reviews.*

[47] Street, S. 2009. In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters. *Philosophical Issues* 19, 273-298.

[48] Strotz, R. 1956. Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies* 23, 165-180.

[49] Suhler, C. and Callender, C. 2012. Thank Goodness That Argument Is Over: Explaining the Temporal Value Asymmetry. *Philosophers' Imprint* 12, 1–16

[50] Sullivan, M. 2018. *Time Biases: A Theory of Rational Planning and Personal Persistence*. Oxford: Oxford University Press.

[51] Thaler, R. 1981. Some Empirical Evidence on Dynamic Inconsistency. *Economics Letters* 8, 201–207.

[52] Tversky, A. and Kahneman, D. 1981. The Framing of Decisions and the Psychology of Choice. *Science* 211(4481), 453–458.

[53] Urminsky, O. and Zauberman, G. 2016. The Psychology of Intertemporal Preferences. In G. Keren and G. Wu (eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making*, Chichester, West Sussex: John Wiley and Sons.

[54] Velleman, J. D. 1991. Well-being and Time, *Pacific Philosophical Quarterly* 72, 48–77.

[55] Weitzman, M. L. 2001. Gamma Discounting. American Economic Review 91(1), 260-271.

[56] Żuradzki, T. 2016. Time-biases and Rationality: The Philosophical Perspectives on Empirical Research about Time Preferences. In Stelmach, J., Brożek, B. and & K. Łukasz (eds.), *The Emergence of Normative Orders*. Copernicus Press, 149-187.

Appendix: The Logic of Temporal Preferences

Let's see that not caring about tenseless calendar date and not caring about tensed temporal perspective together more or less imply EDU (Callender 2021b). A key axiom of EDU is *stationarity*. When axiomatized, EDU is essentially expected utility theory plus stationarity (Fishburn and Rubinstein 1982). Modifying the terminology of Halevy 2015 to suit our purposes, consider outcomes $x, y \in X$, whose values are real numbers, and $t, t' \in T$, the set of times, such that $0 \leq t, t'$, and delays $\Delta_2, \Delta_1 \geq 0$. Then a set of preferences is stationary if at time $t = \tau$

Stationarity $(x, t + \Delta_1) \sim_\tau (y, t + \Delta_2) \iff (x, t' + \Delta_1) \sim_\tau (y, t' + \Delta_2)$.

where $\sim_\tau$ represents indifference at time $t = \tau$. When an agent with stationary preferences ranks options, her decision depends only on the values of the outcomes (x versus y) and the delay between the two outcomes ($\Delta_2 - \Delta_1$). See Figure 4. An exponential discounter represents someone whose preferences satisfy stationarity. Put in tensed language, right now, at time $t = \tau$, the decision maker has preferences such that *neither the calendar date nor the distance from the present* matter to the decision-maker.

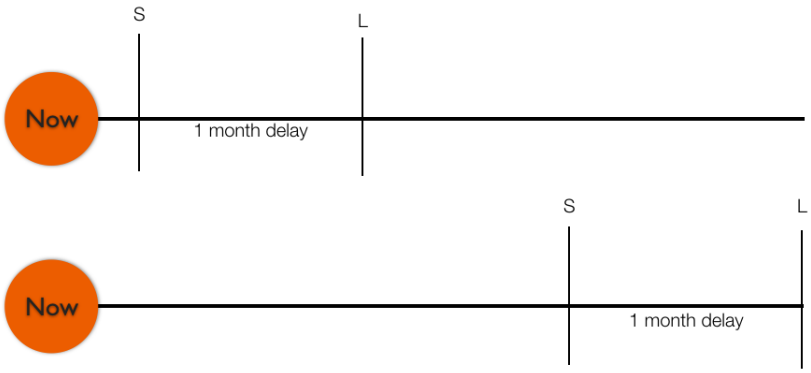

Figure 4: Stationarity: Let the horizontal line represent a tenseless timeline, the dot the evaluation point or Now, S a small reward and L a large reward. A set of preferences that is indifferent between the top and bottom situations is stationary.

Someone might violate stationarity by caring about the calendar date on which the reward was delivered. Add dates to the timeline of Figure 4. Then you can see that such a concern will cause you to violate the condition. Note that although violations of stationary are often associated with preference reversals and inconsistency, that is not accurate. Since the preferences are only elicited at time time $t = \tau$ there is no reversal nor any dynamic inconsistency. At best one might say that a violation of stationarity *sets one up* for inconsistency.

Dynamic consistency requires two evaluation times. We can say a set of preferences at times $t = \tau$ and time $t = \tau'$ satisfies are consistent if

Consistency $(x, t + \Delta_1) \sim_\tau (y, t + \Delta_2) \iff (x, t + \Delta_1) \sim_{\tau'} (y, t + \Delta_2)$.

Consistency looks like stationarity, but note the crucial τ' in the second preference relation. Consistent time preferences mean that one's preferences over temporal outcomes don't change as the present moves from $t = \tau$ to $t = \tau'$, where $\tau' > \tau$. See Figure 5. Suppose that in 2020 one prefers a large later reward in 13 months to a smaller one in 12 months; if time consistent, then one still prefers the larger later reward even when it is only 1 month away compared to the immediate small reward. Violating consistency is to genuinely reverse preferences, which in principle can be exploited. In terms of tense, Consistency can be understood as one's preferences *being insensitive to the tensed "flow of time,"* i.e., your preferences remain the same despite the "orange dot" sliding along the timeline.

A third notion, invariance, acts as a kind of bridge between stationarity and consistency. A set of preferences is time invariant if

Invariance $(x, t + \Delta_1) \sim_\tau (y, t + \Delta_2) \iff (x, t' + \Delta_1) \sim_{\tau'} (y, t' + \Delta_2)$

where $t = \tau$ and time $t' = \tau'$. With invariance, we slide the now *along with everything else*. It tests whether preferences are invariant under a time translation that includes the

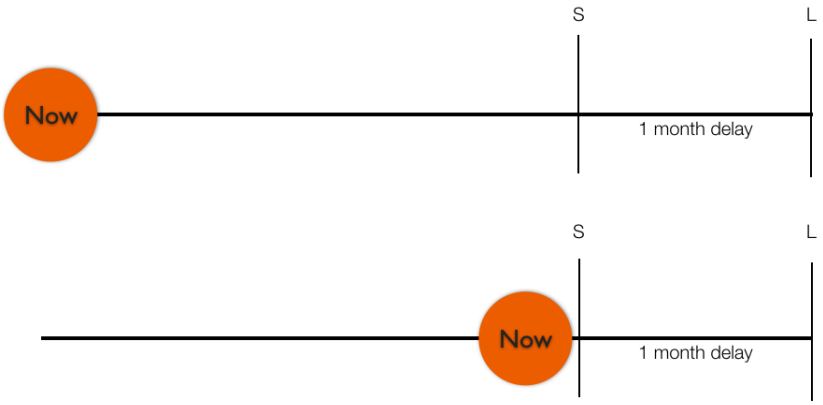

Figure 5: Consistency: Let the horizontal line represent a tenseless timeline, the dot the evaluation point or Now, S a small reward and L a large reward. A set of preferences that is indifferent between the top and bottom situations satisfies consistency.

tensed now. Since the "now" moves, this condition tests whether *any moment of time in the B-series has a special character* or status. See Fig. 6. Invariance is the claim that "preferences are not a function of calendar time" (Halevy 2015, 341). One might violate invariance if he or she cared about particular dates, as in Strotz's example of caring about receiving champagne on the date of his birthday, or simply for some arbitrary reason.

Halevy 2015 states a beautifully simple relationship amongst the three temporal conditions, consistency, stationarity, and invariance, namely:

> Any two implies the third.

The proof is trivial.

This small theorem allows us some insight into what is going on regarding tense. Recall that stationarity more or less implies (if we assume the rest of expected utility theory) exponential discounting. Due to the above theorem, we know that invariance and consistency together imply stationarity. Invariance, we saw, is the condition that preferences are insensitive to calendar time. Consistency meanwhile is the condition that preferences are insensitive to flowing tensed time. Together they imply stationarity, a condition saying that your preferences at a time are sensitive to neither tensed time nor calendar time. As a representation of stationary time preferences, exponential discounting is just that, representing a decision-maker who at a time cares about neither calendar time nor tensed time.

Put the other way around, the theorem implies that a violation of one condition implies a violation of one or both of the others. Hence ¬(stationarity) → ¬(consistency) v ¬(invariance). This means that we might violate stationarity (and EDU) because the tensed flow of time does matter to one's preferences (a violation of consistency) or because the calendar date matters (a violation of Invariance), or both. For the experimental literature on temporal discounting, this permutation is very interesting because

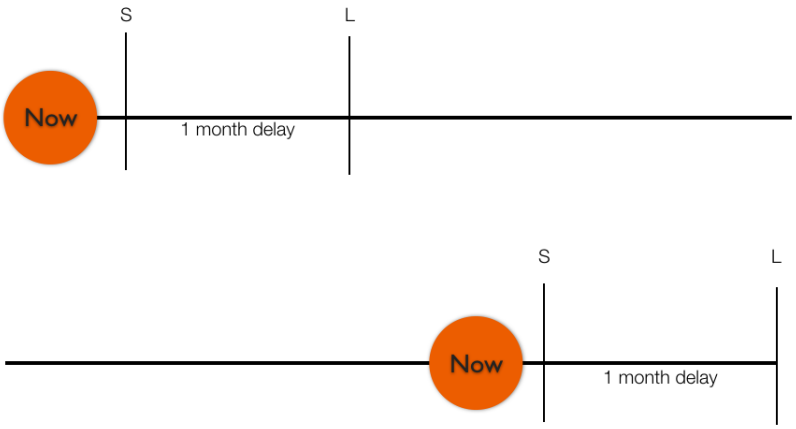

Figure 6: Invariance: Let the horizontal line represent a tenseless timeline, the dot the evaluation point or Now, S a small reward and L a large reward. A set of preferences that is indifferent between the top and bottom situations satisfies invariance.

the smaller-sooner versus larger-later experimental paradigm typically tests at only one time and elicits violations of stationarity. That does not imply a violation of consistency, however, unless we do tests at two times and demonstrate that invariance holds; otherwise the violation may be due to non-invariant preferences. And indeed, the few studies that have tests at two times have found that many decision-makers do violate Invariance (Halevy 2015, Janssens, Kramer, and Swart 2017), thereby showing that dynamical inconsistency cannot be assumed to follow from a failure of stationarity.

Another permutation is interesting: consistency plus calendar time not mattering (invariance) implies that tenses don't matter. This is more or less the argument for temporal neutrality. Temporal neutrality prizes dynamical consistency. Inconsistency threatens utility maximization across a lifetime and can be exploited. And temporal neutralists often ignore "positional" features as impure, such as in the example of Strotz's birthday. Calendar dates are therefore bracketed and invariance assumed (in this regard see Steele 2021). These two assumptions, we now see, are strong enough logically to entail that tenses shouldn't matter.

Finally, we saw that temporal neutrality can be understood in two ways, depending upon what time series we use, a tensed or tenseless one. Invariance states that calendar time doesn't matter. Consistency states that tensed doesn't matter. Suppose we embrace both senses of temporal neutrality. Then both invariance and consistency are assumed to hold; from which, EDU essentially follows.