*Philosophy of Science* (2023), **90**, 1469–1479
doi:10.1017/psa.2023.36

CONTRIBUTED PAPER

# Temporal Neutrality Implies Exponential Temporal Discounting

Craig Callender

Department of Philosophy, University of California San Diego, La Jolla, CA, USA
Email: ccallender@ucsd.edu

(Received 14 January 2023; accepted 31 January 2023; first published online 17 February 2023)


## Abstract

How should one discount utility across time? The conventional wisdom in social science is that one should use an exponential discount function. Such a function is a representation of the axioms that provide a well-defined utility function plus a condition known as stationarity. Yet stationarity doesn't really have much intuitive normative pull on its own. Here I try to cast it in a normative glow by deriving stationarity from two explicitly normative premises, both suggested by the philosophical thesis of temporal neutralism. Putting the argument in this form helps us better understand exponential discounting and challenges to it.


## 1. Introduction

All of our policy decisions have temporally downstream consequences. When building a dam, different choices of material and structure may give it a service life of 50 or 100 years. When making the decision, should we allow our time preferences to affect our choice, and if so, how? Time preferences are preferences for when outcomes are delivered. A very common one is that often people prefer positive outcomes to be delivered sooner than later. Used throughout public policy, discounted utility theory is an extension of expected utility theory that permits time preferences like these. But it is controversial what type of discount function to use.

The founders of discounted utility, Paul Samuelson and Frank Ramsey, introduced a discount function that subtracts utility from outcomes at a constant rate per unit time. In continuous time this schedule of discounting leads to a function with an exponential form. Both explicitly rejected the idea of interpreting this form normatively (Samuelson 1937, 161; Ramsey 1928, 543). Yet within a few decades exponential discounting was widely accepted as the only rational way to discount.

What happened? Strotz (1955), and Koopmans (1960) and Lancaster (1963), demonstrated that an exponential function dominates other functions in the sense that a decision maker using an exponential discount function cannot suffer a preference reversal solely due to their time preferences. Any other function leads to potential preference reversals and hence possible exploitation, a cardinal sin in economics and rational choice theory. The threat of exploitation put exponential discounting on a normative pedestal. As Loewe summarizes,

> After Strotz' contribution, the choice of exponential discounting was not an arbitrary choice anymore, nor a choice of convenience; exponential discounting was found to be now the rational standard in intertemporal choice, one based on the fundamental intuition that any normal person is in fact able to plan ahead. (Loewe 2006, 204)

The result of this normative canonization is that the exponential discount function is widely employed in public policy, from decisions on climate policy to choices of dam materials.

Given its widespread use, it's important to scrutinize the case for treating exponential discounting normatively. Put in axiomatic form, exponential discounted utility becomes essentially a representation of the axioms that give us expected utility theory plus a condition known as Stationarity (Fishburn and Rubinstein 1982). Psychologists and behavioral economists have for a long time reported that people tend not to satisfy Stationarity (Thaler 1981; Loewenstein and Prelec 1992; Urminsky and Zauberman 2016). Is Stationarity nonetheless normatively required? It is widely thought that violations of Stationarity lead to preference reversal and therefore that people tend to discount sub-optimally. This interpretation isn't quite right, however, for Stationarity (as we'll see) concerns preferences at only one evaluation point in time. Violating this condition is not a preference *reversal*, as a reversal is a dynamic process that takes time. In recognition of this fact, a violation of Stationarity is sometimes awkwardly dubbed a "static reversal."

Minus the direct connection to preference reversal, Stationarity loses its normative grounding, and with it, so does exponential discounting. If Stationarity had independent normative purchase on us, this might not be a problem. Yet Stationarity doesn't really have much intuitive normative pull on its own. Compare with the axioms of expected utility theory. These have all been contested in one way or another but most initially strike us as compelling. Not so with Stationarity.

With that brief set-up I can now state the aim of this paper: I offer a new argument for the normative status of Stationarity. The argument takes as its foundation the philosophical thesis known as *temporal neutrality*. I derive Stationarity from a sharp form of this thesis. Crucial to the argument is distinguishing two forms of temporal neutrality and noticing what is needed to derive Stationarity. One can certainly contest the resulting argument. In fact, I will in section 6. Nonetheless, the argument is valid, novel, and based upon independently accepted normative premises; more than that, I feel that it captures the "spirit" behind the imposition of Stationarity. And seeing a clean presentation of the argument better allows us to understand what might make it objectionable.

## 2. Exponential discounted utility and rationality

Discounted utility theory considers a decision maker who must choose at some time $t = \tau$ from among various paths of consumption. These consumption paths are streams of temporally indexed goods. Perhaps one is choosing between apples today

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

and oranges tomorrow versus oranges today and apples tomorrow. Let the vector $x_t = \langle x_{t1}, x_{t2}, \ldots, x_{tn} \rangle$ represent the amounts of the $n$ instantaneous goods to be consumed at time $t$. The decision maker wants to maximize her utility function over these vectors of goods, $u(x_t)$. Because she has time preferences, and in particular, discounts the value of temporally distant outcomes, she modifies her utility function with a discount function, $D_\tau(t-\tau)$. This function represents how she would discount at the decision time, $t=\tau$, and it measures temporal distances from the time of evaluation, i.e., by the delay $t-\tau$. The upshot is that the decision maker aims to maximize

$$u_\tau(x) = \sum_{t=\tau}^{t=\infty} D_\tau(t-\tau)u(x_t), \tag{1}$$

where we assume that $D_\tau(0) = 1$ (that we don't discount the present) and that $0 < D_t \leq 1$ (that the discount function discounts).

To get from discounted utility to the exponential model, one must choose an exponential form for $D$:

$$D_\tau(t-\tau) = \left(\frac{1}{1+\rho}\right)^{t-\tau}, \tag{2}$$

where $\rho$ is the so-called discount rate. The continuous-time version of (2) is $e^{-\rho(t-\tau)}$. Exponential discounting is special in that it is constant through time. It takes the same proportion away from utility in each time period. Because an exponential discounter removes the same amount from utility proportionate to the amount of temporal distance elapsed, *when* the evaluation moment happens is irrelevant to such a discounter. Whether the present is today, tomorrow, or next year doesn't matter, which is why many presentations of exponential discounting often leave the evaluation time $\tau$ out of the formula.

As mentioned, neither Samuelson nor Ramsey endowed exponential discounting with normative significance. What did that was the result by Strotz. 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). He proves that the exponential discount function (2) is the unique function that will lead to time-consistent choices. Of course, one need not discount, and that is consistent with this result because not discounting is constant discounting with $\rho = 0$.

In the well-known Fishburn and Rubinstein (1982) system one essentially derives exponential discounting from five axioms. The first four are commonly employed in obtaining a well-defined utility function. So for present purposes all the action concerns the fifth, Stationarity. 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 $\Delta_2, \Delta_1 0$. Then a set of preferences is Stationary if at time $t = \tau$ they satisfy

$$\textbf{Stationarity} \quad (x, t+\Delta_1) \sim_\tau \big(y, t+\Delta_2\big) \Leftrightarrow (x, t'+\Delta_1) \sim_\tau \big(y, t'+\Delta_2\big).$$

When an agent with stationary preferences ranks options, her decision depends only on two differences, the difference between the values of the outcomes ($x$ versus $y$) and

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

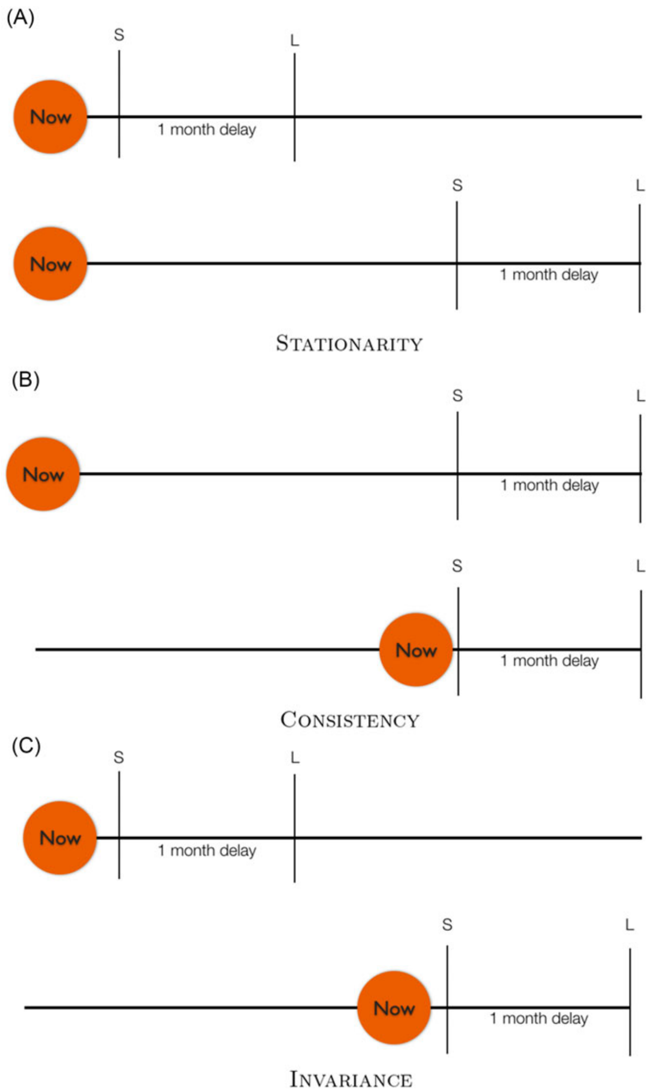


**Figure 1.** Let the horizontal line represent a tenseless timeline, the dot the evaluation point or Now, $S$ a small reward, and $L$ a large reward. A set of preferences meets the relevant condition when the decision maker is indifferent between the top and bottom situations.

the temporal difference or delay between the two outcomes ($\Delta_2 - \Delta_1$). See figure 1(A). *When* the outcomes occur is irrelevant to the exponential discounter.

Although it has been widely known for many decades that Stationarity is not descriptively adequate, thanks to Strotz's result the condition is interpreted as the best way to discount. For instance, textbooks in behavioral economics refer to the axioms of Fishburn and Rubinstein, Stationarity included, as (e.g.) the "axioms of rationality for time discounting" (Dhami 2016, 593).

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

## 3. Problems

The justification for a normative understanding of exponential discounting faces at least two immediate and linked problems. One, a decision maker who violates Stationarity does not necessarily exhibit dynamic temporal inconsistency. And two, Stationarity by itself doesn't seem to have strong normative pull. So, absent the connection to dynamic inconsistency, Stationarity—and exponential discounting—seems normatively unmoored.

The first problem is simple enough to see. To manifest Stationary preferences a decision maker must have two preferences, but these preferences are elicited at the same time, $t = \tau$. There is no reversal. Reversals require two times. Minus a reversal, there is no automatic path to exploitation.

To appreciate the point, compare Stationarity to another temporal condition, Consistency. We can say a set of preferences at times $t = \tau$ and $t = \tau'$ satisfies Consistency if

$$\textbf{Consistency} \quad (x, t + \Delta_1) \sim_\tau \big(y, t + \Delta_2\big) \Leftrightarrow (x, t + \Delta_1) \sim_{\tau'} \big(y, t + \Delta_2\big).$$

Consistency looks like Stationarity, but note the crucial $\tau'$ 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 1(B). Someone who violates Consistency genuinely reverses preferences. In principle that reversal can be exploited. In terms of the figure, the decision maker's preferences change as the "dot"—the now—slides along the timeline. Consider the experimental paradigm typically used in testing our temporal preferences, the "smaller-sooner larger-later" paradigm. At time $t = \tau$ a subject is asked to decide between a small immediate award of \$100 and a larger award of \$120 a week later. They are also asked to decide between the smaller award and larger award but pushed out a year away and a year and a week away, respectively. Studies show that many of us display diminishing impatience (Thaler 1981). We take the small immediate award in the first choice but are willing to wait the week for the larger reward if it is a year away. These non-Stationary preferences are not compatible with exponential discounting. To an exponential discounter, a week is a week and \$20 is \$20, no matter when these occur. But note that having non-Stationary preferences is not enough to be exploited. Suppose someone has the above preference pattern. So long as she sticks to her guns she cannot be exploited. She said she would wait the extra week for the larger reward, and if she still prefers that later, she is not exploitable.

What theorists are implicitly assuming is that she will *not* stick to her guns, that when the smaller reward draws close she will not want to wait the extra week. Of course, we do not usually know that because few experiments test the subjects a second time. In fact, in the few recent experiments that have been done that do ask subjects to return to answer more questions, it turns out that many that do switch do not violate Stationarity (about which more below).

The condition that theorists and experimentalists are implicitly assuming is called Invariance by Halevy (2015). Invariance acts as a kind of bridge between Stationarity and Consistency. A set of preferences is Invariant if

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

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

where $t = \tau$ and $t' = \tau'$. See figure 1(C). With Invariance, we slide the evaluation point *along with everything else*. It tests whether preferences are indifferent under a time translation that includes the evaluation time. Because the evaluation time moves with the rewards, Invariance tells us whether the decision maker cares about some particular events along the timeline. See figure 1(C) and note that the "dot" is the same distance from the reward outcomes in both cases.

In Invariance we have isolated what is needed to connect Stationarity to something normatively charged. That is because Stationarity plus Invariance together imply Consistency. In fact, any two of the conditions imply the third (Halevy 2015). We can prove this using the pictures in figure 1. Note that condition half overlaps with each of the others. Assume that preferences are transitive. Then joining two of the pictures via the common overlap will produce the third picture, for any two pictures; e.g., sliding Stationarity over the overlap of Invariance results in our picture of Consistency. Minus Invariance, however, we lack a path from violating Stationarity to being possibly exploited. And since Invariance is not satisfied in a substantial number of subjects when it has been tested (Halevy 2015; Janssens et al. 2017), it is not merely a technical axiom that can be assumed for the sake of convenience. It is a substantial assumption.

This gap between violations of Stationarity and genuine preference reversals leads to the second problem. Minus the connection to Inconsistency, Stationarity on its own just doesn't seem to have much to recommend it normatively. If Stationarity were independently compelling, we could acknowledge the above gap and simply assume that rationality demands it nonetheless. Stationarity states that if I prefer one temporal stream of outcomes to another, say {eat fish, eat veggies, eat fish} to {eat veggies, eat fish, eat veggies}, then I should also prefer, for any *x*, {*x*, eat fish, eat veggies, eat fish} to {*x*, eat veggies, eat fish, eat veggies}. More aggregate good is supposed to be better. Stationarity assumes that trade-offs in a time period don't affect overall aggregate goodness. Yet holding that hardly seems a dictate of reason.

With the link to preference reversals shown to be incomplete, and with little independent and transparent rationale, Stationarity becomes normatively unmoored. Can we put it on more secure normative footing? In what follows I show that if we understand the philosophical thesis of temporal neutrality a certain way, then Stationarity follows as a deductive consequence of temporal neutrality. To see this, we need to distinguish between two kinds of temporal neutrality, both of which are crucial to the argument.

## 4. Temporal neutrality and tense

The philosophical position known as temporal neutrality can be traced back to ancient times. It is typically associated with Spinoza ([1992] 1687), Adam Smith (1790), Henry Sidgwick (1874), and John Rawls (1971). David Brink provides a recent succinct statement:

> [T]emporal 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

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

> location has no independent significance or no significance per se. (Brink 2011, 358)

Location in time of outcomes isn't by itself a relevant factor when acting rationally. As one can see from the link to rationality, temporal neutrality is an explicitly normative thesis. It is a claim about how best to promote one's well being.

Smith, Sidgwick, Rawls, Brink, and everyone else who writes on temporal neutrality emphasize that the temporal location *can* rationally matter *indirectly*. It is perfectly rational to take the probabilities of outcomes into account when making a decision. Future events are uncertain. If a future outcome depends on a coin flip landing tails, one should take account of the probability of this happening. Time may be a proxy for uncertainty; but ultimately one is then discounting for uncertainty not time. The same goes for concerns about mortality, growth in capital, and much more. Another kind of example is taking calendar date to matter. Strotz (1955) gives the example of wanting champagne delivered on one's birthday. If one orders champagne for their birthday, it makes sense to want it delivered on the day. A late delivery is valued far less than an on-time delivery. Again, a temporal neutralist can endorse this preference, for what has significance is the birthday, not the temporal location itself.

There is an ambiguity in what temporal neutrality means by "temporal location" (Callender 2022). In what type of temporal series is location not supposed to matter? The philosopher John John McTaggart (1908) famously distinguished between two temporal series, an *A-series* and a *B-series*. An A-series organizes events via the temporal predicates {past, present, future}, whereas a B-series organizes events along a timeline ordered by the earlier or later than relation {earlier than, simultaneous with, later than}. In cognitive linguistics, the distinction is sometimes made between *deictic time* and *sequence time*. Deictic time, like the A-series, has an implicit reference to a deictic center, the now, which is often the time of speaker utterance. Sequence time, like the B-series, is simply calendar or clock time, moments related by a directed ordering relation and typically endowed with a metric that provides a measure of duration. The B-series makes no reference to a now. Both A- and B-discriminations are temporal relations, but A-predicates relate an event to a now, a deictic center, whereas the B-relation refers to two explicitly identified events. Because what is the deictic center changes, statements with A-predicates change their truth value depending upon when they are said, unlike statements with B-predicates.

Disambiguated, we can distinguish two senses of temporal neutrality:

> *Tensed temporal neutrality*: temporal location in an A-series should have no significance.
>
> *Tenseless temporal neutrality*: temporal location in a B-series should have no significance.

Tensed temporal neutralism holds that temporal perspective, whether an event is past, present, or future, shouldn't matter to you. "When" you are on your timeline shouldn't count in how you value an outcome. If you think of tenses as a kind of temporal indexical, then the idea is that one shouldn't discount for indexical features. That the time is *now* in addition to being (say) noon, GMT, January 1, 2022, shouldn't matter. In philosophy, tensed temporal neutralism is challenged by cases described by

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

Parfit (1984) who shows that we often have a strong desire to discount outcomes when they go past. Regarding the type of discounting considered here, future discounting, there are few challenges. Temporal neutralists are against any kind of intensification of value to an outcome due to its being near the present.

Tenseless temporal neutralists hold that calendar or clock time shouldn't matter. This kind of discounting isn't much discussed in normative theory, and when it is, it is a bit tricky to define. Earlier I gave Strotz's example of wanting champagne delivered on his birthday. Champagne delivered afterwards has less value. Here, the position in the B-series matters, the birthday. No one thinks that discounting the value of champagne when it is delivered late is irrational. Unqualified tenseless temporal neutralism has no advocates. Recall the way temporal neutralism was described by its advocates: it always includes "per se" type clauses as in the Brink quote. To find a defensible tenseless neutralism, one must isolate away all of these "impurities" like Strotz's birthday. Suppose we moved his birthday too. Then should he care about that position in time? Arguably not (see Sullivan 2018).

Advocates of temporal neutrality are not always clear about what kind of time series they mean. I think the tensed reading captures what most care about. Sidgwick's concern is to counsel people to not give in to impulsive acts that satisfy the momentary preferences of the present self. It advocates for the importance of now-for-later sacrifices, not 2027-for-2032 sacrifices understood tenselessly.

In sum, we've identified two forms of temporal neutralism, each of which has some normative force. The tensed variety has a long history of distinguished champions; the tenseless variety hasn't been noticed as much, but to the extent it has it also has defenders.

## 5. The derivation of exponential discounted utility

Recall that the conditions Invariance and Consistency together imply Stationarity. Stationarity in turn implies (with the usual Fishburn and Rubinstein axioms) the exponential form of the discount function. Our derivation is now very simple. It consists of simply noting that tensed temporal neutralism implies Consistency and that tenseless temporal neutralism implies Invariance.

Look again at Consistency. See figure 1(B). Consistency says that you are indifferent between outcomes that differ only in where the dot is. The dot represents your tensed location, the evaluation point $t = \tau$. As "you" change, you still have the same preferences. If you preferred smaller sooner when the now was earlier, you still prefer smaller sooner when the now is later. In other words, Consistency is simply the expression of tensed perspective—location in the A-series {past, present, future}—not mattering to your preferences, which is the very definition of tensed temporal neutralism.

Turn now to Invariance. The shift from the upper preference to lower preference in figure 1(C) moves the now *and* the temporal location of the reward outcomes (maintaining the same delay from now). The only thing that isn't changed between the upper and lower conditions is the timeline itself. Invariance states that "preferences are not a function of calendar time" (Halevy 2015, 341). In other words, they are not a function of location in the B-series {earlier than, simultaneous with, later than}, which is the very definition of tenseless temporal neutralism.

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

To my knowledge, the connections between Consistency and tense and Invariance and tenselessness haven't been noticed before. Elsewhere I show how this observation provides insight into the "exponential versus hyperbolic" narrative in behavioral economics. Here I simply wish to point out that temporal neutralism, if understood as the conjunction of its tensed and tenseless forms, implies Stationarity, and therefore, the exponential form of the discount function. *When* outcomes occur is irrelevant to the exponential discounter, so someone who doesn't care about A-series position or B-series position will have Stationary preferences.

What is attractive about this motivation of exponential discounting is that it relies on normative principles antecedently accepted by many philosophers. As we saw, Stationarity on its own seemed to have a weak normative basis. Now we can view it as a result of the twin claims that temporal perspective and calendar time shouldn't matter to one's preferences. Of course, the implication is a deduction, so if Stationarity has poor normative standing then so does one or the other of the premises. I'm not suggesting that the implication logically strengthens Stationarity. What I am saying is that Stationarity's normative claim on us before was unclear and that now it's easier to see how it follows from clear normative principles.

## 6. Discussion

Providing a clear argument for exponential discounting from explicitly normative premises assists us in evaluating the standard model. Here I will not assess the argument in full but note some possible replies and developments for further investigation.

### *6.1. The first premise: Consistency*

Consistency says that tensed perspective shouldn't matter. This premise is challenged by what Pettigrew (2020) calls the *problem of changing selves*. Consistency says that as the self—the little "dot" in the figures—moves with time, it continues to honor its previous preferences. As you develop through time, your preferences stay the same. If you picked smaller-sooner, then later you still pick smaller-sooner. Yet of course preferences can change. *You* can change. I once preferred chocolate ice cream to coffee-flavored ice cream; now I don't, and having previously preferred chocolate isn't a mark in its favor now. Have I done anything irrational? In other cases like Ulysses falling under the spell of the Sirens the question gets trickier. There are many responses to the problem of changing selves and I can't do justice to them here (see Hedden 2015; Pettigrew 2020). Here I will simply note that some responses will not insist on Consistency. Our interpretation may need to be restricted to cases where the decision maker's preferences are stable.

### *6.2. The second premise: Invariance*

Our worry about Consistency is that as *you* change *you* may not share the preferences you once had. Invariance says that your preferences should remain invariant as *the world* changes and *you* stay the same. On its face, this condition isn't remotely plausible. Of course calendar time matters. There are all those anniversaries, financial collapses, and wars to take into account.

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

I see essentially two ways to respond.

One is to dig in one's heels and try to tease apart pure from impure temporal preferences, as discussed, and restrict the argument to pure time preferences. I find that very unpromising (see Callender 2021; Ziff 1990). But another response (Steele 2021) holds that (in my language) time preferences can be usefully modeled as independent of utilities when those time preferences are tensed but not when they are tenseless. Tenseless temporal preferences get absorbed into the utility of an outcome. I like this response because it puts the focus on the real worry, namely, preferences that flip flop with tensed perspective. Unlike B-time, A-time contains an essentially indexical component to it. Intuitively, why should we change our preferences—rationally—just because the occupant of an indexical changes?

### 6.3. Getting to zero

There is a long history in philosophy and early neoclassical economics insisting that one should never discount *at all* purely for reasons of time (Peart 2000; Żuradzki 2016). Translated into the current model, this tradition advocates for $\rho = 0$. Since not discounting is constant discounting and constant discounting is exponential discounting ($e^0 = 1$), the present argument is *compatible with this advice* but *doesn't imply it*. If we want to imply no discounting, we might adopt what I'll call Strict Temporal Neutrality:

$$\textbf{Strict} - \textbf{TN} \quad (x, t) \sim_\tau (y, t) \Leftrightarrow (x, t') \sim_{\tau'} (y, t').$$

It states that if you're indifferent now ($t = \tau$) between $x$ and $y$ then you should be indifferent at any other evaluation point ($t = \tau'$) too. The temporal relationship between $x$ and $y$ doesn't matter. Strict-TN implies that $\rho = 0$. Since the delivery times of outcomes do not matter, there is no room for any non-trivial discounting consistent with this condition. Economics assumes that non-trivial discounting is necessary (not wanting future outcomes to swamp current ones), so I didn't consider this form of temporal neutrality; but it is certainly worth examining.

## 7. Conclusion

Exponential discounting is widely viewed throughout social science and policy as the correct way to discount the future. Yet its normative basis has always been a bit shaky. In the preceding I shine a light on Stationarity that casts it in a normative glow. By observing that temporal neutrality comes in two forms, tensed and tenseless, and seeing that these forms imply Consistency and Invariance, respectively, we can derive Stationarity from independently accepted normative premises. This argument also helps us isolate different kinds of challenges to the standard model.

## References

Brink, David. 2011. "Prospects for Temporal Neutrality." In *The Oxford Handbook of Philosophy of Time*, edited by Craig Callender, 353–81. Oxford: Oxford University Press.

Callender, Craig. 2022. "Is Discounting for Tense Rational?" In *Temporal Asymmetries in Philosophy and Psychology*, edited by Christoph Hoerl, Teresa McCormack, and Alison Fernandes. Oxford: Oxford University Press.

https://doi.org/10.1017/psa.2023.36 Published online by Cambridge University Press

Callender, Craig. 2021. "The Normative Standard for Future Discounting." *Australasian Philosophical Review* 5 (3): 227–53.

Dhami, Sanjit. 2016. *The Foundations of Behavioral Economic Analysis*. Oxford: Oxford University Press.

Fishburn, Peter C. and Ariel Rubinstein. 1982. "Time Preference." *International Economic Review* 23 (3): 677–94.

Halevy, Yoram. 2015. "Time Consistency: Stationarity and Time Invariance." *Econometrica* 83 (1):335–52.

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

Janssens, Wendy, Berber Kramer, and Lisette Swart. 2017. "Be Patient When Measuring Hyperbolic Discounting: Stationarity, Time Consistency and Time Invariance in a Field Experiment." *Journal of Development Economics* 126 (C):77–90.

Koopmans, Tjalling C. 1960. "Stationary Ordinal Utility and Impatience." *Econometrica* 28 (2):287–309.

Lancaster, Kelvin. 1963. "An Axiomatic Theory of Consumer Time Preference." *International Economic Review* 4 (2):221–31.

Loewe, Germán. 2006. "The Development of a Theory of Rational Intertemporal Choice." *Revista de Sociologia* 80:195–221.

Loewenstein, George and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and Interpretation." *Quarterly Journal of Economics* 107 (2):573–97.

McTaggart, John E. 1908. "The Unreality of Time." *Mind* 17 (68):457–74.

Parfit, Derek. 1984. *Reasons and Persons*, vol. II. Oxford: Clarendon Press.

Peart, Sandra. 2000. "Irrationality and Intertemporal Choice in Early Neoclassical Thought." *Canadian Journal of Economics* 33:175–89.

Pettigrew, Richard. 2020. *Choosing for Changing Selves*. Oxford: Oxford University Press.

Ramsey, Frank P. 1928. "A Mathematical Theory of Saving." *Economic Journal* 38 (152):543–9.

Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Samuelson, Paul A. 1937. "A Note on Measurement of Utility." *Review of Economic Studies* 4 (2):155–61.

Sidgwick, Henry. 1874. *The Methods of Ethics*. London: MacMillan.

Smith, Adam. 1790. *The Theory of Moral Sentiments*. Oxford: Oxford University Press.

Spinoza, Baruch. (1687) 1992. *The Ethics, Treatise on the Emendation of the Intellect, and Selected Letters*, trans. Samuel Shirley. Indianapolis: Hackett.

Steele, K. 2021. "Why Time Discounting Should be Exponential: A Reply to Callender." *Australasian Philosophical Review* 5 (3):284–95.

Strotz, Robert H. 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23 (3):165–80.

Sullivan, Megan. 2018. *Time Biases: A Theory of Rational Planning and Personal Persistence*. Oxford: Oxford University Press.

Thaler, Richard. 1981. "Some Empirical Evidence on Dynamic Inconsistency." *Economics Letters* 8 (3):201–7.

Urminsky, Oleg and Gal Zauberman, 2016. "The Psychology of Intertemporal Preferences." In *The Wiley Blackwell Handbook of Judgment and Decision Making*, edited by Gideon Keren and George Wu. Chichester: John Wiley and Sons.

Ziff, Paul. 1990. "Time Preference." *Dialectica* 44 (1/2):43–54.

Żuradzki, Tomasz. 2016. "Time-Biases and Rationality: The Philosophical Perspectives on Empirical Research about Time Preferences." In *The Emergence of Normative Orders*, edited by Jerzy Stelmach, Bartosz Brozek, and Lukasz Kurek, 149–87. Göttingen: Copernicus Press.


**Cite this article:** Callender, Craig. 2023. "Temporal Neutrality Implies Exponential Temporal Discounting." *Philosophy of Science* 90 (5):1469–1479. https://doi.org/10.1017/psa.2023.36