Katie Steele¹

Why Time Discounting Should Be Exponential: A Reply to Callender²

According to Craig Callender (2020), the “received view” across the social sciences is that, when it comes to time and preference, only *exponential* time discounting is rational. Callender argues that this view is false, even pernicious. Here I endorse what I take to be Callender’s main argument, but only insofar as the received view is understood in a particular way. I go on to propose a different way of understanding the received view that makes it true. In short: When time discounting is suitably conceived, the exponential form of the discounting function is indeed uniquely rational.

¹ School of Philosophy, Australian National University, Canberra, ACT 2600, Australia. *Email:* katie.steele@anu.edu.au

² Financial support by Riksbankens Jubileumsfond (the Swedish Foundation for Humanities and Social Sciences) is gratefully acknowledged.

1. Time's significance

The temporal position of goods matters to their value in a variety of ways. Consider the following stories about Margot and Amisha, which illustrate two ways in which the timing of eating oysters can matter.

Summertime Margot. Margot loves oysters in the summertime when the air is warm and smelling of salt, and when she is on holidays, in a relaxed frame of mind. It is now the last week of her summer holidays. Oysters are worth a lot to her this week. But she will not sign up for weekly deliveries of oysters, as oysters next week, when she returns to her normal work routine, are hardly worth anything to her. Ditto for oysters the week after that.

Impatient Amisha. Amisha also loves oysters. She loves them all year round. But she is impatient. It is now the last week of her summer holidays. Oysters are worth a lot to her this week. But she will not sign up for weekly deliveries of oysters, as oysters next week, given the wait, are not worth so much to her. Oysters the week after that, given the even longer wait, are worth proportionately less again.

There are many other stories that could be told in which the timing of a good matters to its value.

According to the “received view” of time discounting, however, some ways preferences relate to time are rational while others are not. Followed to the letter, the received view regards Amisha’s preferences for oysters as rational but maintains that Margot’s are irrational. Never mind that Amisha is merely impatient, whereas Margot has more substantial reasons for valuing oysters differently depending on their temporal position. What matters, roughly speaking, is that the effect of time distance on Margot’s preferences for oysters is not steady or linear. For Margot, this week is qualitatively different to all other weeks. Amisha’s preferences, on the other hand, correlate with how far in the future her eating oysters occurs. As a result, Amisha can be represented as having a *time discounting function* for oysters that has the approved *exponential* form, whereas Margot cannot be so represented; her time discounting function is non-exponential. This will be explained more fully in section 2 below.

Craig Callender presents a convincing case that it is not just odd but wrong to regard Margot and Amisha as having different status as far as rationality goes. Callender employs both *modus tollens* and *modus ponens* reasoning. That is, he works backwards—“The conclusion of the received view is wrong and so therefore its premises must be faulty”—and he also works forwards—“The premises of the received view are faulty and so the conclusion is not secured”. I will outline the high points of

Callender's arguments in section 3, with reference to the stories of Margot and Amisha.

Like Callender, I regard both Margot and Amisha as rational, and so consider any account of time discounting that suggests otherwise to be mistaken. Moreover, I agree with Callender that one fairly salient way of reconciling the received view with Margot's rationality is unconvincing (section 4.1). I go on to propose, however, a subtly different refinement of the received view that Callender overlooks (section 4.2). It involves clarifying the conditions under which an agent's preferences can be represented in terms of a time discounted utility model. When these conditions are met, the agent is rational if and only if her time discounting function is exponential. Moreover, by way of conclusion (section 5), I suggest that it is really this refined version of the received view that one well known advocate of exponential discounting—the economist Robert Strotz—intended all along.

2. Inferring the discount function

On the “received view”, the way an agent *discounts goods for time* can be inferred from her preferences over outcomes that differ in the temporal position of goods. More on how this works shortly. Let us first back up and put this talk of time, goods and preferences in the context of the standard account of rational choice: expected utility (EU) theory. We will see that EU theory itself leaves open how the timing of goods affects an agent's preferences. The received view about time discounting amounts to a supplement to EU theory.

On EU theory, an agent is rational just in case she can be represented as having a utility function, u , that satisfies the expected utility principle. According to this principle, her utility for risky options is the probability-weighted average of the utility of the possible realisations (or outcomes) of the option. This can be stated as follows (where A is a risky option, P is the agent's subjective probability function and o_i are the possible realisations or outcomes of A):

Expected Utility Principle. $u(A) = \sum_i P(o_i) \times u(o_i)$

The outcomes o_i are supposed to be ways that the world could be; at the limit of precision they are fully detailed world histories. But little is said about the agent's preferences over these world histories except that they are well ordered. Why an agent prefers one outcome or world history to another is not analysed; it is simply a black box.

There have been various suggested supplements to EU theory, however, by way of adding further structure to outcomes and associated rationales for comparing them.

The general move is to depict outcomes as having specified properties or attributes that each contribute some fixed value to the outcome. For instance, so-called *multi-criteria decision analysis* is a family of approaches to distinguishing and aggregating different sources of value in an outcome. Dietrich and List (2013) have finessed this idea in proposing a single, abstract model that accommodates all the ways in which the value of an outcome may be decomposed into the value of its constituent properties.

The economists' depiction of outcomes as temporal streams of consumed goods can be seen as a specific version of these general approaches to structuring outcomes. On this depiction, outcomes are comprised of temporally-indexed bundles of goods; they are *consumption streams*. Koopmans (1960) for instance, represents outcomes as $x = (x_1, x_2, x_3, \dots)$, where the x_t are vectors of goods. That is, the vector $x_t = (x_{t1}, x_{t2}, \dots, x_{tn})$ lists the non-negative amounts of each of the n goods to be consumed at time t . Depending on the structure of the agent's preferences over these consumption streams, her preferences can be represented by a utility model of greater or lesser simplicity.

Early on, pioneering economists like Ramsey (1928) and Samuelson (1937) singled out, for its nice features, the *exponential time discounted utility model* for evaluating consumption streams. On this model, an agent at time τ has utility, u_τ , for a consumption stream, x , as follows (where t is the time variable, $D_\tau(t - \tau)$ is the agent's time discounting function over time distances given by $t - \tau$, and u is her utility function over vectors of goods, x_t):³

Time Discounted Utility. $u_\tau(x) = \sum_{t=\tau}^{\infty} D_\tau(t - \tau) \times u(x_t)$

Moreover, the agent's time discounting function, $D_\tau(t - \tau)$, has the following form (where ρ is the so-called discount rate):⁴

Exponential Time Discounting. $D_\tau(t - \tau) = (\frac{1}{1+\rho})^{t-\tau}$

That is, each unit of time distance from the present, τ , reduces the value of a bundle of goods by a fixed proportion. Koopmans (1960) showed the full set of conditions under which an agent's preferences over consumption streams may be represented in terms of an exponential discounted utility model. But, as Callender points out, the main property that distinguishes exponential from other forms of time discounted utility is Stationarity (on which more below).

³ The expression here allows for an infinite consumption stream that begins at the present time, τ , but it can be adjusted to incorporate past times or to have an upper limit on the end time of the consumption stream.

⁴ The expression is for discrete time. The continuous counterpart is $e^{-\rho(t-\tau)}$.

We need not consider full consumption streams in order to understand the time discounting of Margot and Amisha. A simpler model will suffice in which outcomes can be represented as pairs (x, t) , denoting that some bundle of goods x is consumed at time t .⁵ In fact, we can go yet simpler, and assume that the agent only cares about the consumption of one type of good. (In the case of Margot and Amisha, it is oysters.) In this case, outcomes can be represented as (s, t) , where s stands for the quantity of the good that is consumed and t is the time at which it is consumed.⁶

The property of Stationarity mentioned above can be spelled out for preferences over the simple outcomes (s, t) .⁷ (It is spelled out in terms of the indifference relation from the perspective of time τ , denoted $\approx_\tau$.) Consider outcome quantities s, r whose values are real numbers, and times t, t' , and time distances $\Delta_1, \Delta_2 \geq 0$. The set of preferences satisfies Stationarity if:

Stationarity. $(s, t + \Delta_1) \approx_\tau (r, t + \Delta_2) \Leftrightarrow (s, t' + \Delta_1) \approx_\tau (r, t' + \Delta_2)$

We see that this condition requires that preferences from the perspective of time τ over temporally-indexed goods depends only on the nature of the goods themselves and the time distance between them, not *where* they are in time on a calendar, so to speak.

It is because Amisha's preferences over the outcomes (s, t) satisfy Stationarity while Margot's do not that Amisha can be represented as having an exponential time discounting function with respect to these outcomes, whereas Margot cannot be so represented.

First consider Margot. Let's say we know (filling in the story a little) that, now at time τ , $(n, 0) \succ_\tau (n, 1)$ and $(n, 1) \approx_\tau (n, 2)$ where n is any given number of oysters, and time is measured in terms of weeks, such that $t = 0$ stands for the period up until one week has passed. Margot does not satisfy Stationarity, since she prefers n oysters now to the same number of oysters in one week's time, but this pattern of preference does not hold when the timeline shifts; it does not hold for n oysters next week compared to the same number one week after that. Any time discounted model of Margot's preferences over the outcomes (n, t) will thus not be one in which the time discounting function is exponential. For instance, a candidate model is one

⁵ Lancaster (1963) appeals to outcomes of this form in his investigation of time discounting. As he points out, one can think of these pairs as special cases of consumption streams, where no goods are consumed (or else the status quo is maintained) for all times except the stated time t when bundle x is consumed.

⁶ For various reasons, Fishburn and Rubinstein (1982) and Halevy (2015) use this very simple kind of outcome in their respective investigations of time discounting.

⁷ Like Callender, I appeal to Halevy's (2015) definition of Stationarity (although with some differences in notation); In the next section I will appeal too to Halevy's (2015) definitions of the further conditions of Consistency and Invariance.

where her time discount factor for week 0, $D_\tau(0)$, is one, while her time discount factor for the two subsequent weeks is $w < 1$. That is, $D_\tau(1) = D_\tau(2) = w$. If u_M is Margot's utility function over varying amounts of oysters, her utility at time τ for the outcome $(n, 1)$ is $1 \times u_M(n)$ whereas her utility at time τ for both $(n, 1)$ and $(n, 2)$ is $w \times u_M(n)$.

Now consider Amisha. We know that $(n, 0) \succ (n, 1) \succ (n, 2)$. Specifically, Amisha prefers n oysters during the first week to n oysters during the second to the same extent as she prefers n oysters during the second week to n oysters during the third week. Her preferences satisfy Stationarity: what determines her preference for n oysters is the distance between the temporal locations of these oysters and not where the locations lie on the calendar. So if we were to model Amisha's preferences at the current time τ in terms of a time discounted utility model, her discounting function would be exponential in form. For instance, a candidate model is one where $D_\tau(0) = 1$, $D_\tau(1) = 0.5$ and $D_\tau(2) = 0.25$. In this way $\frac{D_\tau(2)}{D_\tau(1)} = \frac{D_\tau(1)}{D_\tau(0)}$. If u_A is Amisha's utility function over varying amounts of oysters, her utility at time τ for outcome $(n, 0)$ is $1 \times u_A(n)$, her utility at time τ for outcome $(n, 1)$ is $0.5 \times u_A(n)$ and her utility at time τ for outcome $(n, 2)$ is $0.25 \times u_A(n)$.

3. Callender against the received view

On the received view, only exponential time discounting is rational. But what exactly does that mean? Callender suggests it means that agents like Margot (as compared to agents like Amisha) are irrational. In that case, the received view must be something like the following:

Received View. *If an agent's preferences at some time τ can be represented in terms of a time discounted utility model, she is rational if and only if the form of her time discounting function is exponential.*

Whether or not this is really the dominant way of thinking about time discounting is, of course, a sociological matter (hence Callender quotes various economists, psychologists and evolutionary theorists in support of his attribution). Here I will simply take for granted that Callender has accurately represented the views of social scientists on discounting.

Callender goes on to argue against the received view, understood as above (referred to from now on as Received View). For starters, he finds the conclusion that agents like Margot are irrational to be absurd. Moreover, by drawing on an insightful theorem by Halevy (2015), he shows that the conclusion rests on shaky foundations. This section retells, using the stories of Margot and Amisha, Callender's important

criticisms of Received View. This sets the stage for the next section, which explores how Received View may yet be slightly adjusted so that it is defensible.

3.1 Bad conclusion

Let us begin with the absurdity of claiming that Margot (unlike Amisha) is irrational. First, to call someone irrational suggests that they are confused—that their preferences and motivations are in some sense incomprehensible. But that is not true of Margot. Her preferences, since they do not involve mere impatience, are arguably even more comprehensible than Amisha's. More generally, there are a variety of reasons why the temporal position of goods matters. After all, time is correlated with many things—ageing, the seasons, an evolving climate, technological innovation, all sorts of cultural trends—that serve as the background to goods and affect how and the extent to which these goods are enjoyed. It would be very surprising if the impact of these various temporal factors (as we might call them) could in all cases be expressed as exponential discounting of the goods in question. (And more surprising again if the exponential rate of discounting was the same for all goods, as some models suggest.) That is, there is nothing about exponential discounting that makes it the more obvious and explanatory way in which the temporal position of a good affects an agent's preferences.

Moreover, it is not as if exponential rather than non-exponential discounting leads to personal ruin. Callender notes that any suggestion to the contrary is a prime example of normative views clouding one's empirical observations, something that should, as far as possible, be avoided. (He thinks this theory-ladenness has not been avoided, nor well appreciated, in the literature. That is why he refers disparagingly to the received view as “a package”: in addition to the normative claim about exponential discounting, there are psychological and evolutionary findings that purportedly support the normative claim, showing the problems that befall non-exponential discounters and offering complicated explanations for why many people might have this defect.) As Callender points out, both exponential and non-exponential discounters may have discounting curves that lead to later regret, not to mention financial and social poverty. After all, exponential discounting curves can be as steep as you like, effectively recommending that one “live like there is no tomorrow”, placing very little value on goods to be enjoyed at future times. So if due concern for one's future self is the issue, exponential discounting is not the answer. Rather, the answer is no discounting at all, at least for the goods that are crucial to well being at a time.

3.2 Bad premises

So what exactly is the problem with some, but not all, forms of time discounting? As Callender points out, it was Strotz (1956) who made explicit what economists regard to be the real sin: *dynamic inconsistency*. The idea is that one's preferences over goods consumed at a given time should not depend on one's current vantage point. For otherwise, one would be vulnerable to exploitation. For instance, let's say that I currently do not place a high value on having oysters two weeks hence. So I am willing to sell the rights to oysters at that time for a small price. Suppose too that once two weeks have passed, I like oysters a great deal and am willing to pay a lot for them. Someone could make a profit by initially buying my two-weeks-in-the-future oysters at a low price and then later selling them to me at a high price. In general, whenever my preferences change, I can be exploited in this way. If I can predict with certainty that my preferences will change, the exploitative trades or bets can be engineered (a so-called "Dutch book") so that I already know in advance that I will suffer a loss. That seems bad, and a problem for any form of time discounting that permits it.

The property in question (or rather its converse) can be referred to as Consistency, and formally defined as follows (where outcomes are represented as before, t and t' are times where goods are located, $\approx_t$ is indifference from the perspective of time t and $\approx_{t'}$ is indifference from the perspective of time t'):

Consistency. $(s, t) \approx_t (r, t') \Leftrightarrow (s, t) \approx_{t'} (r, t')$

We see that this condition specifies that one's preferences for outcomes do not depend on one's vantage point, or the time at which one has the preferences.

Let us grant that rationality in the dynamic setting requires Consistency.⁸ What Callender draws attention to is that, assuming already a time discounted utility model, the Consistency condition does not entail Stationarity (and with it, exponential time discounting), as many take Strotz's (1956) theorem to show. Nor does Stationarity entail Consistency. This is evident just by looking at the form of the two conditions. Stationarity concerns preferences at a particular time (cash out in terms of the indifference relation $\approx_t$ that is indexed to the time t). Consistency, on

⁸ Callender does not in fact grant this. He cites well-known criticisms of (diachronic, in particular) Dutch book arguments. Some of these criticisms can be disarmed by emphasising that what matters (and what is the target of Consistency) are one's *plans* for whether and how one will change one's preferences over time. These plans are faulty or irrational when one can *anticipate* sure loss. As such, it does not matter whether or not one spontaneously changes one's preferences over time in an unanticipated way. No-one can guarantee a Dutch book against you if you change your preferences in unforeseen ways. So the conviction that one should be able to spontaneously change one's preferences at any given time is not reason to reject diachronic Dutch book arguments, nor, ultimately, Consistency.

the other hand, concerns the relationship between preferences at different times (how $\approx_\tau$ relates to $\approx_{\tau'}$). Something else is needed to bridge the logical gap between Stationarity and Consistency.

The something else is an assumption in Strotz's theorem that seems to have gone unnoticed.⁹ Halevy (2015) draws attention to the suppressed assumption, dubbing it *Invariance*. It can be formally stated as follows (where the notation is as per the statement of the previous conditions):

Invariance. $(s, \tau + \Delta_1) \approx_\tau (r, \tau + \Delta_2) \Leftrightarrow (s, \tau' + \Delta_1) \approx_{\tau'} (r, \tau' + \Delta_2)$

Invariance requires that one's preferences with respect to time have the same form, whatever position on the calendar one currently occupies. That is, what matters is how far away in time the goods are from one's current vantage point.¹⁰ Halevy (2015) goes on to make clear the logical role that Invariance plays. He reveals, in Callender's words (p. 13), "a beautifully simple relationship amongst the three temporal conditions, Consistency, Stationarity, and Invariance, namely: Any two implies the third."

The significance of bringing Invariance to light is that we see that one does not need to be an exponential discounter (satisfying Stationarity) in order to be immune from Dutch books (satisfying Consistency). One just needs to violate Invariance in a way that preserves Consistency. Moreover, this seems a perfectly reasonable approach to time discounting. Indeed, Callender notes that empirical studies suggest that many and perhaps most people who violate Stationarity in fact violate Invariance rather than Consistency.¹¹

Margot is plausibly one such person. As the story goes, the reason Margot loves oysters this week is that she is still on summer holidays and so is in the mood for them. Given her prediction that she will not be in the mood for oysters next week or the week after, when she is back at work, she does not currently value consuming oysters in those weeks. That's as far as our story of *Summertime Margot* goes—it concerns only Margot's preferences at a single time. But the most plausible extension of the story is one where Margot sacrifices Invariance for Consistency.

⁹ Note that this can hardly be regarded the fault of Strotz, who does not try to hide any of the assumptions supporting his (1956) theorem. I return to this point in the concluding section.

¹⁰ This is subtly different to the Stationarity condition, which required that preferences over outcomes depend only on the nature of the outcomes and the time distances between *the outcomes*. Invariance on the other hand says that the preferences over outcomes depend only on the nature of the outcomes and their time distances from the *agent*.

¹¹ Callender cites the empirical results of Halevy (2015) and Janssens et al. (2017).

Consider:

Summers-end Margot. A week has now passed and Margot is back at work. Just as she predicted, she does not particularly value oysters at this time. Nor does she particularly value having oysters next week, since she predicts she will then be similarly blasé about oysters. Ditto for the foreseeable weeks after that while her work routine continues.

(Note that this story and the discussion below are framed in terms of Margot's actual attitudes once she returns to work. But it could be retold in terms of Margot's predictions or *plans* for what her attitudes will be at this later time.¹²) The combination of preferences in *Summertime Margot* and *Summers-end Margot* seem entirely reasonable. They satisfy Consistency, since Margot does not change her preferences for the dated oysters over time. From the start, she did not much value having oysters in either week 1 or week 2, and this attitude remains unchanged as time passes and week 1 becomes the present. Margot thus cannot be Dutch booked. She maintains Consistency, however, by violating Invariance. At the earlier time, she loves oysters *now*, but not so much *later*. The effects of time distance are not invariant, however, with the passing of time. When she returns to work, it is not the case that Margot loves oysters *now*, but not so much *later*. Rather, she neither loves oysters *now* nor *later*.

We see that Margot satisfies Consistency and thus cannot be Dutch booked, but is there nonetheless something bad about her pattern of preferences? Callender argues that Received View can only be saved if Invariance is given some normative defence. He suggests that, at best, Invariance might be defended on the grounds of non-arbitrariness. Roughly, it would be arbitrary and thus irrational for one's time preferences to depend on one's current position in calendar time. But he goes on to say that this line of defence is wanting. Indeed, it is hard to see how such a line could be made convincing in the face of stories like Margot's.

It is worth noting too that "non-arbitrariness" is a rather ambiguous notion in this context. Callender quotes such luminaries as Adam Smith, Henry Sidgwick and John Rawls as advocates of non-arbitrariness of time preferences. But it is a stretch to claim that these scholars argue specifically for non-arbitrariness *in the sense of Invariance*. After all, each of Stationarity, Consistency and Invariance encode a non-arbitrariness *of sorts*. In fact, the quotes Callender appeals to all suggest a stronger form of non-arbitrariness than any of these three conditions. They can be interpreted as prescribing that the temporal position of a good should make no difference

¹² Recall footnote 8.

to an agent's preferences at all. This amounts to no time discounting of goods whatsoever, at any time. Many would think that this is far too restrictive and thus not credible. In any case, it is not a carefully targeted argument for Invariance.

Callender concludes that there is no normative motivation for Invariance, and so there are insufficient grounds to accept Received View. He thus claims that rationality does not require that time discounting be exponential. In fact, he suggests that once the pillar of Invariance is removed, the entire time discounted utility model may come crashing down. Perhaps, after all, it does not make sense to model an agent's outcomes and preferences in a way that makes time an independent, separable factor, especially once we move beyond simple settings involving a single good, as per the stories of Margot and Amisha. Rather, the timing of a good may be better seen as integral to its description, since it may impact on the enjoyment of the good in a way that is idiosyncratic to the good in question.

4. The received view redux

While Callender is right to insist that agents like Margot are not irrational, he is too quick to dismiss the notion that exponential time discounting has a special claim to rationality, not to mention the usefulness of the time discounted utility model. This section explores how the received view may be refined; I argue that it should be adjusted to reflect stricter conditions for when an agent's preferences may be modelled in terms of a time discounted utility model. In the conclusion I will suggest that these were the conditions Strotz originally put forward.

4.1 First pass: focus on *genuine* time discounting

Before getting to the main event, let us consider one salient way of making the conditions stricter for employing a time discounted utility model; a strategy that Callender rightly regards as flawed. This is to insist that a time discounted utility model of an agent's preferences is apt only if it separates out the role truly played by time itself, as opposed to some other properties correlated with time. The idea is that rationality requires only that *genuine* time discounting, in the sense just described, need have exponential form. The view can be spelled out as follows:

Received View⁺. *When an agent's preferences at some time τ are represented by a genuine time discounted utility model, she is rational if and only if the form of her time discounting function is exponential.*

The advantage of Received View+ is that it allows one to say that apparent counterexamples to the irrationality of non-exponential time discounting, like the case of Margot, simply are not genuine cases of time discounting. After all, it is not time itself that makes a difference to Margot's evaluation of oysters. Rather, it is the presence or absence of warm salty air and Margot's levels of stress that make for better and worse oyster eating. So an appropriate model of Margot would require a more complex description of the outcomes. There would need to be at least two types of goods: say, "holiday oysters" and "routine oysters".¹³ Once the goods that matter to Margot are fully accounted for in these terms, we see that there is no further effect, owing purely to temporal position, on Margot's preferences. She does not exhibit any discounting for time whatsoever. (And so she satisfies, in a trivial way, all three conditions of Invariance, Consistency and Stationarity.)

While this may seem a natural enough move in the case of Margot, the idea that we would be able to distinguish what is and what is not purely an effect of time on preferences seems implausible. Even the phenomenon of impatience, which seems to be the paradigmatic effect of time itself on preferences, could be construed in non-temporal terms. What really matters to the agent is not temporal position, but rather the feelings of frustration due to waiting that accompany the good, one might say. In general, it would be very difficult, perhaps impossible, to locate the line between the effects of time on an agent's preferences versus the effects of other properties correlated with time. And even if this line could be located, this would trivialise exponential discounting as a norm of rationality. In many cases, as per Margot, the agent's time discounting would have exponential form, but only trivially, in that she does not discount for time at all. Moreover, she may yet violate Consistency for other reasons.

4.2 The way forward: focus on *whatever counts as time discounting*

Rather than try to pin down the real or *genuine* effects of time on preference, we can focus simply on what is a *useful* characterisation of the effects of time on preference. By way of spelling this out, I draw attention to the different roles that may be played by the Invariance condition. Callender takes Invariance to be an ostensive norm of rationality, albeit not a very convincing one. One way to maintain that it is a convincing norm is to argue, as per above, that apparent counterexamples do not truly undermine the norm—they trade on a faulty depiction of outcomes. But there is another reading altogether of the Invariance condition. As opposed to a norm, it may

¹³ Margot's outcomes would then be represented $\langle s, r \rangle$, where s is the quantity of holiday oysters and r is the quantity of routine oysters.

serve rather as a fixed point or assumption that allows one to see how distance in time affects an agent's preferences. The idea is that once we model an agent as having preferences over outcomes that are Invariant, then we can read off from the model the role that time itself plays in her preferences. After all, the effects of temporal distance, if they really are simply effects of temporal distance, do not depend on the agent's present calendar location. The agent discounts for temporal distance in whatever fashion her Invariant preferences over time suggest.

The model of Margot that we obtain when we *assume* Invariance is the same one described above. But the way we arrive at the model is different. We do not ask what are the properties of outcomes that must be accounted for in describing Margot's preferences that can be distinguished, objectively speaking, from time itself. Rather, we search for a way of describing outcomes so that Margot satisfies Invariance. (There will be such a description of Margot's preferences just in case she is, or plans to be, a stable person over time. Unstable agents who lack any systematicity in their preferences at different times are arguably not very interesting topics of study.) In Margot's case, the obvious way to do this is to discriminate holiday oysters and routine oysters, as before. Then Margot satisfies Invariance in a trivial fashion: the temporal distance of consuming oysters, at any given time, plays no role in her preferences. Her time discounting is thus rational because she satisfies, trivially, Consistency, and therefore Stationarity too.

There are other stories that better demonstrate how this approach—assuming Invariance—is distinct. Consider:

Stockpile Sarita. Sarita loves oysters all year round, but she likes them better the fresher they are. She can store an evening's worth of oysters in her freezer. The freezer keeps the oysters pretty well, but their freshness nonetheless deteriorates linearly with time. As such, Sarita currently values consuming the oysters in proportion to how far away in time this would occur—when she would be removing them from the freezer and eating them. Moreover, she predicts that at any later date, she will similarly value consuming the oysters in proportion to how far away in time this would occur.

If we aimed to disentangle the genuine effect of time distance on Sarita's preferences, we would end up with a trivial model of her time discounting, as per Margot. That is because it is clearly not really temporal distance *per se* that matters to how Sarita values consuming the oysters taken from her freezer. Rather, it is their relative freshness, which is simply a physical attribute of the oysters. We would conclude that the apt model of Sarita's preferences is one that distinguishes all the different grades of oysters in terms of their freshness. There would be no further aspect of Sarita's preferences to be explained by distance in time itself.

One could model Sarita as caring nothing for temporal distance itself at all points in time; this is one way in which she can be modelled as satisfying Invariance. But that is to unnecessarily forgo the benefits of a simple and elegant model of Sarita's preferences. We can just as well model Sarita as having a non-trivial time discounting function. Sarita satisfies Invariance when her outcomes are described simply as a given amount of oysters at some temporal distance, represented (s, t) . Moreover, the story indicates that at all times temporal distance matters to her in a *linear* fashion (satisfying Stationarity). Sarita's preferences can thus be depicted in terms of a time discounted utility model where the time discounting function is exponential. So we see, in a very simple and vivid way, that she is Consistent and therefore arguably rational.

This approach sidesteps unnecessary debates. By way of another example, consider a slight modification of the story *Impatient Amisha*. In this case, Amisha's impatience is such that she greatly values having oysters in the near future, while having them at more distant times has more or less the same low value to her. Let us assume that at future times too, her impatience takes the same form. It would be silly to agonise over what is Amisha's *genuine* relationship with time. Does her impatience mean that time really matters to her? Or is impatience, too, reducible to other properties of outcomes, e.g., feelings of frustration? Consequently, is the right model of Amisha one in which she exhibits non-exponential discounting, or rather one in which she does not discount for time at all? This question is a mere distraction. All that matters is whether Amisha satisfies Consistency. It is clear that she does not, since *there exists* a model of her preferences that satisfies Invariance and yet violates Stationarity (such that her time discounting is non-exponential). Therefore her preferences violate Consistency.

This brings us back to the received view about time discounting. My amended version of it can be stated as follows:

Received View*. *If an agent's preferences at all times τ within a relevant period of time can be represented in terms of a time discounted utility model, she is rational if and only if the form of her time discounting function is exponential.*

For an agent's preferences at different times to be modelled by a time discounted utility model—the requirement in Received View*—she must satisfy Invariance. Arguably, a time discounted utility model is indeed inherently temporally extended in this way. That is, it makes no sense to say that an agent's preferences at a single time can be represented by a time discounted utility model. Only when we consider what is the common form of an agent's preferences at multiple times can we see what role distance in time, rather than calendar time, plays. For a stable agent at

least, distance in time has the same effect, whatever the agent's position in calendar time. And it is the effect of distance in time that is usefully captured by the discounting function in a time discounted utility model. Furthermore, on pain of irrationality, this discounting function must have exponential form.

5. Conclusion

By way of conclusion, I note that Strotz himself proposed something along the lines of Received View*. While he does not explicitly assume Invariance, it is implied by his insistence that a time discounting function capture the role that time distance (and not calendar time) plays in the preferences of an agent. Indeed, Strotz states in the introduction of his paper (1956, 165): "What is crucial to all this is that the discount applied to a future utility should depend on the time-distance from the present date and not upon the calendar date at which it occurs."

Moreover, Strotz takes pains to point out in a long passage (1956, 167–8) that his model can accommodate the effects of calendar time too (such that whatever the nature of an agent's preferences, so long as she has a certain stability, she can be modelled by an Invariant time discounted utility model). Instead of simply fine-graining the description of goods to account for calendar time (e.g. holiday oysters versus routine oysters), Strotz pursues a slightly more elegant model. He makes the agent's timeless utility function for the bundle of goods at some calendar time depend not just on the vector of goods but also on the calendar time. So where in the expression for Time Discounted Utility (refer back to section 2) there appears $u(x_t)$, Strotz has instead $u(x_t, t)$. This allows his model to incorporate calendar time as a separable aspect of a bundle of goods that nonetheless has an effect on these goods that is timeless; it does not depend on distance in time relative to the agent.

Strotz implies that it falls on the modeller to determine a time discounted utility model that is suitably representative of an agent as she moves through time. That is the spirit of Received View*. On this approach, exponential time discounting is, after all, privileged. It is the uniquely rational form of discounting in that exponential discounters alone satisfy Consistency.

References

Callender, C. 2018. 'The Normative Standard for Future Discounting', *The Australasian Philosophical Review*, forthcoming.

- Dietrich, F. and C. List. 2013. 'A reason-based theory of rational choice', *Noûs*, 47(1): 104–134.
- Fishburn, P. and A. Rubinstein. 1982. 'Time Preference', *International Economic Review*, 23: 677–694.
- Halevy, Y. 2015. 'Time Consistency: Stationarity and Time Invariance', *Econometrica*, 83: 335–382.
- 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.
- Koopmans, T. C. 1960. 'Stationary Ordinal Utility and Impatience', *Econometrica*, 28(2): 287–309.
- Lancaster, K. 1963. 'An Axiomatic Theory of Consumer Time Preference', *International Economic Review*, 4(2): 221–231.
- Ramsey, F. P. 1928. 'A Mathematical Theory of Saving', *The Economic Journal*, 38(152): 543–559.
- Samuelson, P. A. 1937. 'A Note on Measurement of Utility', *The Review of Economic Studies*, 4(2): 155–161.
- Strotz, R. H. 1956. 'Myopia and Inconsistency in Dynamic Utility Maximization', *The Review of Economic Studies*, 23(3): 165–180.