5

Intertemporal Choice: Choosing for the Future

Daniel Read, Rebecca McDonald and Lisheng He

5.1 Introduction

Perhaps the majority of our choices are *intertemporal* ones, in which the options under consideration differ in their timing as well as (perhaps) a host of other things. Often these demand we choose between 'smaller sooner' (SS) and 'larger later' (LL) options, in which receiving more or better outcomes comes at a cost in timing. We make such a choice when deciding between immediately going on a short trip or saving our holiday allowance towards a longer but later getaway; or when choosing between getting the current version of a smartphone today or the next version once it is released. More broadly, you might choose to spend your money on consumption today, or invest the money to let yourself consume more later (since your investment will earn you interest). Choices might not even be between similar things: you might deliberate between buying a new TV now or saving towards a new car in a few months' time or having a more comfortable retirement. This choice is still intertemporal, though the options differ in many ways apart from their timing and magnitude[1].

Discussion of intertemporal choice has a long history. In around 450 BC, Plato provided a surprisingly modern account in his dialogue *Protagoras*. He suggested (through his spokesperson Socrates) that outcomes further away in time are diminished in perception just as are objects farther away in *space*. Plato considered this a bias, asserting that the correct way to choose between was to ignore outcome delay and focus on outcome size, just as we should ignore distance when estimating the heights of distant objects:

> Do not objects of the same size appear larger to your sight when near, and smaller when at a distance? … And the same holds of thickness and number; also sounds, which are in themselves equal, are greater when near, and lesser when at a distance … Now suppose doing well to consist in doing or choosing the greater, and in not doing or avoiding the less, what would be the saving

[1] Strictly speaking, intertemporal choice is not always between SS and LL. For example, people might choose between having a child now (putting their career on hold) and waiting to have a child later (prioritising their career in the meantime). This is more likely to be viewed as a struggle about when is the best time to schedule a good thing. Later we will take a broader view of intertemporal choice.

> principal of human life? Would it be the art of measuring or the power of appearance? (Plato, translated 1956, Section 356d).

Plato's *normative* account is based on the view that it is rational to take the larger of two outcomes, regardless of when they occur.

> If anyone says … the pleasure of the moment differs widely from future pleasure and pain, to that I should reply: And do they differ in anything but in pleasure and pain? There is nothing else … If you weigh pleasures against pleasures, you of course should take the more and greater; or if you weigh pains against pains, you should take the fewer and the less; or if pleasures against pains, then that course of action should be taken in which the painful is exceeded by the pleasant, whether the distant by the near or the near by the distant; and you should avoid that course of action in which the pleasant is exceeded by the painful (Plato, translated 1956, Section 357b).

However, Plato's *psychological* account is one in which people, if insufficiently trained, are governed by their current time perspective and prefer sooner things ('seeing' them as if they were closer) rather than later. The issues Plato identified, involving the conflict between normative and descriptive accounts of decision-making, remain a central part of current discussions of intertemporal choice.

Many terms have been used for the driving force behind intertemporal choice, including delay of gratification, delay discounting, time discounting, impatience and time preference. All these terms capture, in a general sense, the rate at which people are willing to exchange earlier for later consumption. In economics, this is called the marginal rate of intertemporal substitution. The core concept is that outcomes are intrinsically less important to us the later they occur.

Before continuing, try making an intertemporal choice of your own. Imagine you are offered a choice between $100 today, or $120 in one year. You can be sure because of the integrity of those making the offer (i.e., the authors of this chapter) that you will receive the chosen sum at the chosen time. Would you take *SS* ($100 today) or *LL* ($120 in one year)? Exactly how much would the later amount have to be for you to be just willing to wait? Put your answer in the blank:

I would be indifferent between $100 today and $_______ in one year.

We will refer to the number you gave as X. By telling us X, it should be as though you were telling us, 'For any amount less than $X, I would choose *SS* over *LL*, but for any amount more than $X, I would choose *LL* over *SS*.'

Using X, you can now work out your personal discount rate for money[2]. Set up a fraction with X as the numerator and the earlier amount of $100 as the denominator, and subtract one from that fraction. For example, if you demanded $110 a year, then your discount rate would be 10 per cent, or 110/100 – 1 (in percentage terms). Our guess is that many readers will have wanted even more than the $120, implying a discount rate of at least 20 per cent.

5.2 The Normative Model

Let us analyse the choice between $100 now and $120 in one year. According to conventional economic analysis, there are three influences on whether you would take the $120. These are (1) the amount of interest the $100 today could have earned during the year; (2) the amount of interest it would cost you to borrow $100 and pay it back in one year; and (3) *the experimental interest rate*,

[2] Strictly speaking, this is a personal interest rate r, or the rate at which money must grow to make you indifferent between now and later. The term *discount rate* is sometimes reserved for a simple transformation of this interest rate, $r/(1+r)$, but in this chapter (following convention), we will use *discount rate* to refer to both the personal interest rate and the personal discount rate.

which is given by (X/S)–1, in this case, 20per cent[3]. To see why these are the only things that matter, imagine your bank offers a 5 per cent rate of return on savings, and charges a 25 per cent rate of interest on borrowing. (The savings rate of 5 per cent is a bit high by current standards, but the 25 per cent is about what it costs to borrow using a credit card.) If you take the $100 now, you can save it and then have $105 to spend in one year's time. So you should not specify X less than $105. In general, you should never take LL over SS if by doing so you would earn less than your best alternative (in this case, taking S = $100 and saving it in the bank). What if the experimental interest rate was higher than your borrowing rate in the market? Then you should always choose LL. You should reject our $100 today, borrow $100 from the bank to spend instead and then in one year pay back the bank loan (plus the 25 per cent interest) with X, making a profit from the difference between X and the amount you owe the bank. If the experimental interest rate is less than 5 per cent or more than 25 per cent, the decision is easy. When you specified X, you *should not* have given a value less than the amount you would earn by investing the money, or greater than the

[3] The general formula for this rate is given by $\left(\frac{L}{S}\right)^{1/t} - 1$ where L and S are, respectively, the larger and smaller amount, and t is the interval separating them. In our example, 1/t = 1, so it drops out.

amount it would cost you to borrow the money, regardless of how patient or impatient you are about consumption now and later.

What if the experimental interest rate is in between 5 per cent and 25 per cent, as it is when X is $120? Then what you should choose depends. Mostly it depends on your personal financial circumstances. If you are currently saving money at 5 per cent, and you can get a return of 20 per cent on some money, you should take LL. The reason is that even if you want to spend the money right now, you would be better off withdrawing $100 from your savings at the bank and spending that, and then depositing the $120 in the bank in one year. You will then be $15 better off, having spent $100 now, than if you had taken SS from the experimenter. If, on the other hand, you are currently in debt and paying 25 per cent interest, then you should take the $100 now and reduce your debt by that amount. That way you will pay $5 less in interest over the coming year.

But suppose you have no savings and no debt? This situation is a bit more complicated, because it depends on your consumption preferences. The basic intuition is clear. Roughly, if you would prefer to consume $100 now over $120 in one year, you should take SS; otherwise, you should take LL. More precisely, what you should choose depends on which option brings you to the highest indifference curve (discussed next). Since the person with no savings and debt can achieve different distributions of consumption by taking SS or LL, it is possible for either option to lead to the better distribution for them. A thorough discussion of these issues can be found in Cubitt and Read (2007). It should be emphasised, however, that this intermediate case will only apply to those people who are not currently saving or holding debt, and who are offered a choice implying an interest rate between their best saving and borrowing opportunity.

In reality, people do not have perfect access to saving and borrowing opportunities, so the choice of X might violate the economic predictions to some degree. To illustrate, if you are offered $100 in cash now as opposed to a direct deposit of $120 in one year, and you want to spend the $100, and would have to walk to an ATM to withdraw that amount, you might be willing to forego some future money to save yourself the walk. Or if you had to fill out a form and wait to borrow $100, you might forego some interest to not have to fill out the form. But these violations are likely to be quite minor, so the scale of violations typically observed in experiments and surveys (and, perhaps, in your own introspection when you previously specified X) is difficult to reconcile with the economic model. This is particularly true since the overwhelming majority of these violations imply discount rates that are *too high* (see Frederick, Loewenstein and O'Donoghue, 2002). Explaining this observation is one of the key objectives of behavioural economics. For instance, many people will prefer SS even when the experimental interest rate is greater than the rate of interest on their credit cards.

Let us take a step up in generality and consider a framework first formalised by Irving Fisher (1930) in the *Theory of Interest*. Imagine a person with a stock of resources. If they used up all their resources today, they would enjoy consumption at an amount C_1. Alternatively, those resources could all be saved up for tomorrow. Assume a *single* external rate for saving and borrowing, *r*, and that there are no additional costs to saving and borrowing. So if all of the resources were saved up, consumption next period could be $C_2 = C_1(1 + r)$. Alternatively, consumption can be spread out between the two periods. This gives an *intertemporal budget constraint* that describes all possible ways to distribute consumption over the two periods, given that any consumption not taken in Period 1 grows at interest rate *r* until it is consumed in Period 2.

The optimal point on the intertemporal budget constraint depends on the person's indifference curves for consumption across the two periods. The indifference curve plots all equally desirable combinations of C_1 and C_2, so the slope of the curve reflects their marginal rate of substitution (i.e., the rate they are willing to trade off) between consumption at times 1 and 2. For simplicity, let us first suppose a decision maker has a stable discount rate r^*, regardless of the allocation to each period. If her personal discount rate is equal to the external discount rate (i.e., $r^* = r$), her indifference curves are straight and parallel with the budget constraint line. Thus they are equally happy at any point along the budget constraint. If $r^* > r$, she will allocate all to C_1, and if $r^* < r$, she will allocate all to C_2. However, people usually prefer to smooth their consumption, and thus their indifference curve will be convex to the origin (i.e., $C_1 = C_2 = 0$), because when their consumption is concentrated in one period, they will forego quite a lot of it in that period for a small increase in consumption in the other period. As usual for optimisation problems, the highest attainable indifference curve is the one that is at a tangent to the budget constraint. Figure 5.1 illustrates such a case. This person will choose to allocate his or her consumption at point A.

Figure 5.1.

How consumption is allocated across two time periods.

Fisher (1930) assumed each person's time preference is determined by personal characteristics.[4] He also assumed that actual behaviour (intertemporal choices) would be influenced by each individual's current consumption circumstances and the individual's expectation of future consumption circumstances (since this would influence that person's perceived endowment of

[4] Fisher identified these as foresight, self-control, habit, life expectancy, concern for the lives of others and fashion. This list covers many of the most prominent elements of behavioural economic models of intertemporal choice.

consumption in later periods, and hence his or her budget constraint). For example, a university student with a low current income might expect to earn more after graduation. She may therefore ‘borrow’ from her future self to smooth her consumption over time. Her choices might appear to give more weight to present than future consumption and be interpreted as showing a high level of impatience. Later, the same person might enjoy a good salary, but expect to retire in the next few years, when her income would be much lower. She may then reduce her consumption in the present and reallocate it to her future, retired self. This behaviour looks patient, as though she gives future consumption more weight in her decisions than present consumption. But conceivably, her private discount rate over this lifetime might be unchanged (she may be as impatient for dessert at age fifty as she was at age twenty). This demonstrates a nuance of intertemporal choice that is often overlooked in simple models of time preference: even if a person has positive, stable underlying time preferences, she may act in ways that appear (to a naïve observer) to be inconsistent so as to smooth her lifetime consumption.[5]

[5] According to Fisher (1930), pure time preference is the rate of substitution when the allocations to the two periods are equal (i.e., the slope of the indifference curve on points $C_1 = C_2$ in Figure 5.1). [[Compositor: Please hyperlink this figure reference to Figure 5.1.]] Pure time preference defined in this way is often not observable and is different from measured time preference elicited in experiments. Thus, if not specially stated, time preference (or discount

5.3 The Core Findings in Intertemporal Choice Research

The rational theory of intertemporal choice just laid out has important implications. One is that an individual's personal degree of impatience will not influence the intertemporal decisions he or she makes for money, at least not for relatively small amounts. We cannot describe this implication more concisely than Richard Thaler (1981):

> In the case of perfect capital markets, everyone behaves the same way *at the margin* since firms and individuals borrow or lend until their marginal rate of substitution between consumption today and consumption tomorrow is equal to the interest rate (p. 201, italics added).

As a corollary to this, each person's discount rate will be approximately the same in all settings involving monetary choice. The best known findings in intertemporal choice are called 'anomalies' because they are inconsistent with this economic model, or close relatives of that model similarly designed to reflect

rate) in this chapter refers to the measured rate of time preference (or discount rate).

economic rationality (see Thaler, 1981; Loewenstein and Thaler, 1989; Roelofsma, 1996).

Most anomalies concern how differences in options influence preferences in ways that generate deviations from the predictions of a model based on a constant discount rate. There are two classes of anomalies: in the first class, differences in the objective properties of options influence preferences over those options; in the second class, economically irrelevant variations in choice circumstances or option description produce variations in preference. Here we give a brief tour of both classes of anomaly. In all cases, we will be discussing the *discount rate,* defined as the experimental interest rate which would render the decision-maker indifferent between an option now and an option later. For example, as discussed previously, someone indifferent between $100 now and $120 in one year will be said to display a 20 per cent discount rate for that decision.

5.3.1 The Sign Effect

The *sign effect* is that the discount rate is lower when choosing among delayed losses than delayed gains (Thaler, 1981). For instance, if we compare two choices, between *receiving* $100 today or $120 in one year, and between *paying* $100 today or $120 in one year, we would find a much higher proportion would pay today than would receive in one year. This contradicts the view that a single rate governs all decisions since anyone with a discount rate of less than 20 per cent would take the later money in the first choice and pay today in the second, while anyone with a rate greater than 20 per cent would show the reverse preference. The sign effect is even more extreme than this, since many people will prefer to pay now rather than to pay later, even if there is no financial advantage in waiting and indeed even if the amount to be paid is *greater* now than later (e.g., Hardisty, Appelt and Weber, 2013; Yates and Watts, 1975). Hardisty and Weber (2009) found that some domain differences in intertemporal choice (e.g., quantities of health are discounted at a different rate than quantities of money) are partly attributable to the differences in outcome signs, with health changes typically described as losses, and monetary changes described as gains.

5.3.2 The Magnitude Effect

The *magnitude effect* is that the discount rate is higher when choosing among small options than larger ones (Thaler, 1981). For instance, if we compare choices between receiving $100 today or $120 in one year, or receiving $1,000 today or $1,200 in one year, we would find a much higher proportion would choose to receive the larger delayed outcome than the smaller one. Again, this contradicts the idea of a single rate that governs all intertemporal decisions.

5.3.3 The Delay Effect

The *delay effect* is that people require lower rates of interest for long delays than for short ones (Thaler, 1981). For example, someone indifferent between $1,000 today and $1,200 in one year (implying a required interest rate of 20% per annum) will prefer $1,440 in two years to $1,000 today (implying a required interest rate lower than 20 per cent per annum). Although the delay effect is often interpreted as evidence for hyperbolic discounting, it confounds hyperbolic discounting with subadditive discounting (Read, 2001), which is illustrated in the following section.

5.3.4 Subadditive Discounting

Subadditive discounting occurs when people are more impatient when the delay to an outcome is divided into subintervals than when it is undivided (e.g., Kinari, Ohtake and Tsutsui, 2009; McAlvanah, 2010; Read, 2001; Read and Roelofsma, 2003; Scholten and Read, 2006; 2010). Reconsider the example for the delay effect discussed in the preceding section. The preference for $1,440 in two years over $1,000 today could be due to a lower required interest rate for the interval between one year later and two years later (i.e., hyperbolic discounting) or a lower required interest rate for the two-year delay when it is undivided, relative to when it is divided into two one-year intervals (i.e., subadditive discounting).

5.3.5 The Common Difference Effect (and Its Reversal)

The *common difference effect* is that discounting over an interval changes as a function of the onset of the interval (Loewenstein and Prelec, 1992), with less discounting the later the interval begins. Tests of the common difference effect serve as evidence for hyperbolic discounting, free from the confusion of subadditive discounting. However, existing evidence has shown a mixture of constant impatience (e.g., Halevy, 2015; Read, 2001), decreasing impatience or hyperbolic discounting (e.g., Bleichrodt, Gao and Rohde, 2016; Green, Myerson and Macaux, 2005; Kinari et al., 2009; Read and Read, 2004; Read and Roelofsma, 2003; Scholten and Read, 2006; Sopher and Sheth, 2006) and increasing impatience or antihyperbolic discounting (e.g., Attema et al., 2010, 2016; Holcomb and Nelson, 1992; Read, Olivola and Hardisty, in press; Sayman and Öncüler, 2009).

5.3.6 Sequence Effects

Intertemporal choices between sequences of positive outcomes have been shown to be different from those between single-dated outcomes. Most compellingly, *negative time preference*, which means a preference for a positive outcome to take place later rather than earlier, has been frequently observed when subjects choose between intertemporal sequences (Loewenstein and Prelec, 1993; Loewenstein and Sicherman, 1991; Read and Powell, 2002; Rubinstein, 2003), but is rarely observed when they are asked to choose between two single-dated outcomes. However, this preference for increasing sequences is not at all a universal phenomenon (Gigliotti and Sopher, 1997; Manzini, Mariotti and Mittone, 2010; Scholten, Read and Sanborn, 2016). More research is needed to determine when people prefer increasing sequences, and when they prefer to take rewards early.

5.3.7 Framing Effects Versus Procedural Invariance

Normatively equivalent elicitation methods should elicit the same preferences for the same choice options (Tversky, Sattath and Slovic, 1988). However, diverse framing effects found in intertemporal choice have demonstrated violations of this principle of *procedural invariance*. The *delay/speed-up* *asymmetry* is that people tend to be more impatient when an intertemporal choice is described as SS being delayed to LL than when it is described as LL being expedited to SS (Loewenstein, 1988; Scholten and Read, 2013; Weber et al., 2007). The *date/delay effect* is that people tend to be less impatient when time is described with calendar dates than with the length of delays (LeBoeuf, 2006; Read et al., 2005). *Outcome framing effects* are that people exhibit different degrees of impatience when the outcomes are framed in different terms, such as an interest rate, the gross interest earned or the total amount earned (Read et al., 2005; Read, Frederick and Scholten, 2013). The *(asymmetric) hidden-zero effect* is that the explicit display of getting nothing at a later time in the SS option reduces impatience (e.g., Magen, Dweck and Gross, 2008; Read et al., in press; Wu and He, 2012). This research is in line with an argument put forward by Ebert and Prelec (2007) that the impact of time is uniquely 'fragile', meaning that small manipulations can markedly increase the impact of time on choice, or reduce that impact almost to nothing.

5.3.8 A Brief Note on Neuroeconomics

The wide range of anomalies just described illustrate just how much even monetary discounting (which theoretically should be constant for all outcomes and roughly constant across people) is highly malleable, and contingent on a wide range of cognitive and affective processes (see Lempert and Phelps, 2016, for a further review). An important approach to understanding intertemporal choice is through the relatively new field of neuroeconomics (Loewenstein, Rick and Cohen, 2008). This approach attempts to look directly at processes occurring in the brain to pinpoint physiological drivers of intertemporal choice. Much neuroeconomics has focussed on whether intertemporal choice is governed by dual systems (Faralla et al., 2015; McClure et al., 2004, 2007;) or a unitary system (Kable and Glimcher, 2007, 2010). A second line of studies tries to reveal the neural underpinnings that determine individual differences in impatience by comparing brain-lesion patients with healthy people (Manuck et al., 2003; Sellitto, Ciaramelli and di Pellegrino, 2010), associating activation in certain brain regions with impatience measures from intertemporal choice (Hariri et al., 2006; Kayser et al., 2012; Smith et al., 2015) or external control of certain brain regions (Figner et al., 2010; Kayser et al., 2012). The neural evidence is growing rapidly but yet remains inconclusive (see Kable, 2013, for a review). It is likely, however, that in future versions of this chapter the classic anomalies in intertemporal choice will be further clarified by greater understanding of the physiological processes that lead to choice in intertemporal settings.

5.4 Functional Forms and Inconsistency

In the previous section, we discussed empirical anomalies to the Fisher model. Here we return to the idealised economic models and go more deeply into the problem of dynamic inconsistency. The original, convenient and elegant model of time preference is exponential, formalised by Paul Samuelson in 1937. This model assumes that time preference can be captured by a single parameter, the discount rate. This is applied through an exponential function, so consumption worth X after a delay of t years is worth $\frac{X}{(1+r)^t}$ today with an annual discount rate of r. This model is analytically tractable and also provides a normative framework, since it will lead to consistency over time in decision-making ('dynamic consistency'). To illustrate dynamic consistency, imagine a person today who has a 10 per cent discount rate, applied exponentially. She will be indifferent between receiving consumption worth 110 utils[6] in one year or 121 utils in two years, since her present value of both 110 utils in one year and 121 utils in two years is precisely the same at 100 utils. Now imagine one year passes. Will she still be indifferent between the options she had considered the

[6] A *util* is a measure of the benefit or utility from consumption. A related term in the literature is 'instantaneous' utility, referring to the utility experienced at a given moment, which could be expressed in utils.

previous year? The 110 utils are now available. The 121 are now due in one year, and given exponential discounting also have a present value of 110 utils. So her relative present value of the two outcomes is unchanged by the passage of time, and her choices will be dynamically consistent.

We discussed preferences over utils here, rather than over amounts of money, but experiments often involve people choosing between money. To interpret discount rates elicited from money amounts as revealing and underlying discount rate over consumption (or even utils), these studies make two assumptions. The first is that there is a direct link between monetary experimental receipts and consumption (or utils), so that £1 today gives £1 of consumption (or its equivalent utils) today, while £100 in one year gives £100 of consumption (or its equivalent utils) in one year, or at least it gives anticipated or perceived consumption of those utils at those times. Second, they assume 'narrow bracketing' in which respondents think about the choices in isolation, ignoring the opportunity costs of their decisions as determined by financial markets, as discussed in the earlier section.

The exponential discounting model and the dynamic consistency it implies is not, and indeed was never expected to be, descriptively accurate. To see how, let us return to *Protagoras*. In it, Plato describes a person who follows 'the power of appearance' rather than 'the art of measuring'. Plato is not optimistic about how this will turn out. Speaking of the power of appearance, he suggests:

> Is it not that deceiving art which makes us wander up and down and take at one time the things of which we repent at another, both in our actions and in our choice of things great and small? But the art of measurement would invalidate the power of appearance and, showing the truth, would fain teach the soul at last to find lasting rest in the truth, and would thus save our life (Plato, translated 1956, Section 356d).

We can illustrate with an example of comparing the heights of two skyscrapers, where the taller is slightly farther. When both are distant, the taller correctly appears to be taller. However, if you walk towards the skyscrapers, eventually the relative heights of the skyscrapers will reverse in your perception, so the closer but shorter one will appear to be largest, until the farther but taller one disappears. This is a powerful metaphor for time inconsistency and impatience.

While we are in conversation with the ancients, consider Odysseus and the sirens. Odysseus longs to hear the sirens, but their song is deadly as it lures sailors to crash their ships onto the rocks of their island. It is in Odysseus's long-run interests to hear the song of the sirens without crashing, and we can think of ‘hear the song, don’t crash the ship’, as his original plan[7]. However, Odysseus knows that in the heat of the moment, he would not be able to resist the temptation to run his ship onto the rocks. That is, he cannot trust his short-run self to carry out the original plan. This is the classic failure of dynamic consistency.

This passage from the Odyssey motivates Robert Strotz’s (1955) classic paper ‘Myopia and Dynamic Inconsistency in Dynamic Utility Maximization’. He describes the behaviour he would expect of an agent, like Odysseus, with dynamically inconsistent preferences. He proposes this agent starts life with a stock of resources which must last his entire life. He makes a consumption plan for those resources. The plan involves relatively even spreading, with some privilege given to earlier consumption, so that the consumption plan declines over time. This is the initial plan, which is depicted as the darkest line in Figure 5.2. Strotz suggests that when this agent has the opportunity to reevaluate his plan, he will reallocate his consumption, favouring his (then) current interests over those of his future selves. This allows a little less consumption for later.

[7] Note that it might not have been in his interest to not hear the song and not crash – it is not clear he did his sailors a favour by denying them the opportunity to hear the Sirens even at the cost of their lives.

Subsequent reevaluations will repeat this, with the agent drawing more and more consumption from his future selves into the present.

Figure 5.2.

How consumption is reallocated over time (based on Strotz, 1955).

Strotz proposes that someone who does not recognise this tendency in himself will behave as a 'spendthrift' or else as a 'miser', 'his behaviour being inconsistent with his plans'. O'Donoghue and Rabin (1999) refer to these unaware actors as *naïfs*, contrasted with *sophisticates* who anticipate their potential inconsistency. Sophisticates can adopt one of two broad classes of strategy to alleviate or forestall their inconsistency. The first is the use of consistent planning in which knowledge about future behaviour is taken into account at the time of the initial plan. That is, the initial plan is made such that even the impatient future self will leave the plan unchanged. The second strategy is precommitment. For example, the agent could lock her resources into a pension plan that can only be released at a much later time, so the greedy early selves cannot pilfer it. This second type of strategy is the one chosen by Odysseus, who had his sailors plug their ears with wax so they were not tempted by the Sirens' song. Then he had them bind him to the mast of his ship, ears unplugged, so he could hear the song without being able to default on the original plan (i.e., 'don't crash the ship'). Odysseus, indeed, found a way to cheat the system by having his cake and eating it too, by hearing the Sirens without crashing the ship. Here is a passage from the translation by Lawrence of Arabia (T. E. Lawrence) himself:

> Only myself may listen [to the sirens], after you have so fastened me with tight-drawn cords that I stand immovably secured against the tabernacle of the mast … and if I beg or bid you to let me loose, then must you redoubly firm me into place with yet more bonds (Homer, translated 1932, p. 217).

A little introspection will provide plenty of examples where we might behave like Strotz's agent, or like Odysseus would have if not fastened to the mast. Consider any time you have decided to eat healthily, then given in to a slice of cake, or planned to save carbon by taking the bus, but succumbed to the temptation to drive instead. There are also occasions where we employ a strategy of consistent planning: we might stuff ourselves with healthy snacks so we do not want any cake, or plan our travels for times when the roads will be so busy we will not want to drive. We might instead precommit, by giving the cake away to our colleagues, or buy a week's bus pass in advance and hand our car keys over to a friend.

Given the power of the metaphor and the abundance of real-life examples, it is not surprising that situations where we might expect violations of the predictions of the exponential model have been extensively studied. As mentioned earlier, this has often been done using SS-LL monetary amounts, although these studies have not been so conclusive. In nonmonetary domains, however, evidence for Odysseus-like violations are widespread as in the earlier-mentioned papers by Shapiro (2005) and Read and Van Leeuwen (1998).

Here we briefly outline some models of intertemporal choice that attempt to account for this type of behaviour. Hyperbolic discounting models, first named by Chung and Herrnstein (1961) and strongly advocated by Ainslie (1975, 2001), are the best established class of models to account for Strotz-like preference reversals (for a review, see Andersen et al., 2014). In these models, the discount rate declines with the time horizon between now and the event, approaching zero as that horizon becomes very remote. The most familiar variant of hyperbolic discounting is the one proposed by Mazur (1987), where consumption worth X today is worth $\frac{X}{1+kt}$ after a delay of time t. This has two major consequences. First, when we look into the future and compare two distant outcomes with one more delayed to the other, these outcomes are discounted to a very similar extent, and so we may choose the larger later outcome. However, the same difference separating outcomes close to us in time generates a much bigger difference between the present values, so we might wish to choose the smaller sooner outcome.

An alternative form, which has become an extremely influential model in economics, is quasihyperbolic (or, ‘β-δ’) discounting (O’Donoghue and Rabin, 2000; Phelps and Pollak, 1968; Elster, 1979; Laibson, 1997). Additional weight is placed on the current period, but the rest of the time horizon is characterised by exponential discounting. This allows preference reversals as proposed by Strotz, but is mathematically convenient. The format for this model is exponential discounting, but with an additional weight $\beta < 1$ that applies to every period after the current one, giving a ‘present bias’. Formally:

$$\frac{X}{(1+r)^t} = X \qquad \text{if } t=0$$

$$\frac{\beta X}{(1+r)^t} \qquad \text{if } t>0$$

Quasihyperbolic discounting differs from ‘true’ hyperbolic discounting in that it does not lead people to discount far-future outcomes at a very low marginal rate. To illustrate, imagine someone evaluating two future outcomes – the loss of one thousand lives in twenty years, or the loss of one thousand lives in forty years. For a quasihyperbolic discounter, the proportional reduction in value for each additional year would be the same regardless of when that year occurs (as long as it was not the first, which would be discounted more). For a true hyperbolic discounter, however, the proportional reduction in value is less for each year of additional delay. True hyperbolic discounting entails that people will care more than they 'should' about outcomes delayed for a very long time, as well as less than they 'should' about outcomes delayed for a very short time.

It is often claimed that the delay until an outcome drives a decline in the discount rate, which motivates the hyperbolic models we discussed. However, there is an alternative interpretation. Declining discount rates are observed in SS-LL studies where SS is now and LL is increasingly later. But this design confounds the delay from now until LL with the interval between S and L. What if the interval was driving the declining rates, and not the delay? This proposition motivated the development of the subadditive discounting model (Read, 2001).

The subadditive models assumes that the per-period discount rate elicited from SS-LL choices when a long interval separates them is lower than the per-period discount rates elicited from the same interval when it is subdivided. The discounting function is then given by $\frac{X}{(1+r)^{(t_L - t_S)^{\psi}}}$, where ψ captures the nonlinear perception of time (individuals are assumed to perceive short intervals as longer, and long ones as shorter, as explained in Read, 2001). The difference $t_L - t_S$ captures the delay until the later and earlier outcome (and hence the interval between them) and r is, as usual, the discount rate. The subadditive model is consistent with a wide range of empirical results (Glimcher and Kable, 2007; Kable and Glimcher, 2010; McAlvanah, 2010; Read, 2001; Read and Roelofsma, 2003), although it may be that both subadditive and hyperbolic discounting work together (see Scholten and Read, 2006). The idea that distortions in time perception drive intertemporal choice (at least in part) is now well established (Takahashi, 2005; Zauberman et al., 2009).

These are by no means the only models proposed to explain behaviour that cannot be captured in an exponential discounting framework (see the review in Doyle, 2013). Here, we have highlighted the role for nonexponential models of discounting that may apply across domains and in many real-world applications. It is worth noting that the majority of existing nonexponential models are designed to accommodate relatively few effects and mostly focus on explaining the purported tendency to prefer SS in choices between now and later, but LL when the choices are between later and even later (i.e., the common difference effect). However, as described previously, there is an abundance of evidence for other violations of the normative model.

So far, we have considered discounting for money, consumption or utility. Under standard economic theory, if commodities, experiences and outcomes can be translated into one another (i.e., if they are fully fungible), then there is no need to look any further than discounting for utility, since agents will distribute their resources into a consumption stream that maximises discounted utility. However, it is an open empirical question whether a single-time preference rate for each person is sufficient to explain behaviour across all contexts, or whether discounting could be domain-specific. We take up this question next.

5.5 Commodity Specific Discounting

To think about discounting for different commodities (e.g., candy and money), we need information about how someone would be willing to trade across commodities and across time. We first need to distinguish between the own-rate and money-rate of discount. The own-rate is the discount rate of a good measured in terms of itself. For instance, if Heather thinks six candies one year from now are as good as three now, her candy-rate of interest is 100 per cent. The money rate of discount depends on what Heather is willing to pay for candies. Suppose she is indifferent between $1.00 worth (to her) of candies today and $1.50 worth of candies in one year. Her money rate of discounting for candies is 50 per cent. For this analysis, we do not mind how many candies this equates to. However, to see whether her money rate and her own-rate of discounting are consistent, we need to add this information.

Suppose Heather is willing to pay $1.00 for three candies today. If she is consistent and rational (in economic terms), we can make four inferences. First, since she is indifferent between three candies today and six the next year, and indifferent between three candies today and $1 today, she will also be indifferent between $1.50 next year and six candies next year. She will be willing to pay $1 today for six candies next year and will be willing to pay $1.50 in one year to get three candies *today*. Finally, she will treat paying $1 today as equivalent to paying $1.50 in one year, so her discount rate for money is the same as her money discount rate for candies. For Heather, all four outcomes are exactly as good as one another.

A similar analysis can be carried out when there is market for a good. We can analyse that market to determine what the commodity's own-rate and its money-rate should be. You can think of this as a scaled-up version of Heather's preferences, with the difference that these are now market prices as opposed to her private willingness to pay – the equivalences are between market prices and quantities of a commodity. If the market perfectly clears, it will imply consistency between the outcomes over time. Like for Heather, in the market, the four outcomes are precisely equivalent in their value. Figure 5.3 shows forward interest rates (the price to transform outcomes at one time to outcomes at a different time) and spot rates (the price to transform outcomes of one kind into outcomes of another kind).

Figure 5.3.

Spot and forward rates for candies and money showing equivalence between outcomes.

Why do the own-rates of interest differ from commodity to commodity? Keynes (1936) described three determinants of discount rates: The *yield*, the *carrying cost* and the *liquidity premium*. The yield is the benefit from using a good. For some goods, earlier receipt is better because you can start using them right away and get more total use out of them. This can apply to durables such as houses, land and radios, but usually not to consumables such as candies that are likely to be consumed at the same rate regardless of when they are received and will then be largely forgotten. *Carrying cost* refers to the costs of having to hold on to something for a period of time, and includes things such as storage, obsolescence, spoilage and wastage. An increased carrying cost will decrease the discount rate because if you want to consume something in the future, you prefer not to have it until then – it is better not to receive tomorrow's dinner today. The third determinant is the *liquidity premium*. This refers to how easily something can be converted into something else. Cash money is the most liquid, since it can be easily converted into other types of consumption. But if that money is invested in a pension, it cannot easily be converted into other consumption (at least, not until you retire). Keeping money liquid, therefore, is worth something.

A separate issue, which might make a larger amount proportionally less attractive than a smaller, is *satiation*[8]. For example, imagine a choice between two donuts today and four in one year. Suppose you take the two donuts today. This does not mean you discount the pleasure-from-donuts at a rate of 100 per cent or more (as one might naïvely assume from your choice of SS). Instead, the later donuts are less valuable because of your *diminishing marginal utility for donuts* as well as your *donut discount rate*. These two reasons can be tricky to disentangle empirically, although recent studies have tried to address this by either double elicitation of both utility curvature and the discounting function (Andersen et al., 2008) or avoiding utility curvature through their experimental design (Attema et al., 2010; Laury, McInnes and Swarthout, 2012).

5.6 Approaches to Measuring Commodity-Specific Discounting

We have discussed various reasons, compatible with rational decision-making, that could generate different strengths of preference for sooner and later receipt across different goods even if the underlying discount rates are identical. This

[8] Keynes was concerned with aggregate demand, not individual preferences. In the aggregate, satiation is a less important issue since it takes a lot of a commodity to satiate everybody.

line of reasoning is often not pursued by researchers investigating the discounting of different goods, and yet these issues pervade all types of intertemporal choices. Indeed, a large number of results show domain-specific discounting can be easily accounted for by Keynesian factors and satiation. Further, after taking account of diminishing marginal utility, liquidity, carrying costs and yields, the normative models have at their heart the assumption that there is one 'true' underlying discount rate. This can be thought of as the discount rate for utility over time (we are using 'utility' in the sense of a measure or index of good that can be added up, as in Edgeworth (1879), though we recognise there are many different interpretations of the concept)[9]. However, many scholars take the view there are intrinsic differences between the discount rate for different commodities. For example, Berns, Laibson and Loewenstein (2007) suggest a key role for the determinants of intertemporal choice that are unrelated to discounting per se. They focus on anticipation, self-control and representation and argue for a unifying theory of intertemporal choice that accounts for these effects. But to explore these issues and develop such a theory, we need a way to compare discount rates across different commodities. We can conceive of four possibilities.

[9] See Kahneman, Wakker and Sarin (1997) and Read (2007) for reviews of the concept of utility in the context of modern behavioural economics.

The first possibility is to examine the amount of a good that, if received today, is equivalent to receiving some other amount at a later date (Option 1). This is the essence of the SS-LL experimental paradigm, but is vulnerable to all the practical constraints we have discussed so far: it requires a proportional relationship between quantities of the good and their value and perfect divisibility of goods. This often renders the approach unsuitable. For example, remember Heather, who is indifferent between three candies today or six candies in one year. Imagine she is also indifferent between two pens today or three pens in one year. If we can say for sure that six candies are valued twice as highly as three candies, or that three pens are 1.5 times as good as two pens, then we can infer that her discount rate for candies is 100 per cent while that for pens is 50 per cent. But we cannot be sure about this relationship. A more fundamental question is whether there is any meaningful comparison between the 100 per cent for pens and the fifty for candies. If these numbers represent the proportional loss of 'utility' that arises from a one-year delay to each of them, they would be comparable, and this is the approach often taken in the literature (e.g., Chapman, 1996, elicits proportional loss for health and for money and directly compares these). But it is an open question whether these numbers can really be compared directly. Another serious problem with this approach is that it restricts us to looking at divisible goods, and ones for which we can assume constant marginal utility (or else correct for nonlinearity).

Since the first possibility – eliciting own-rate discounting – seems problematic, we next consider approaches that share the same foundation: making use of some common currency or 'numeraire' to allow comparison of discounting across commodities.

The first of these is to find the 'numeraire equivalent' value of an outcome occurring today (Option 2). We can then ask, if I delay the outcome by *t* days, what is the 'numeraire equivalent' of this delayed outcome? That is, how much of the numeraire today would be equivalent to receiving the good after the delay? By taking the proportion of the numeraire that is lost when delaying different commodities, we have elicited the *numeraire-rate* of discounting for each of them. These discount rates are directly comparable to one another, although the rate of discount for different numeraires will not necessarily be the same[10]. Crucially, however, the ordering of the magnitude of discount rates should be the same regardless of the numeraire. The alternative numeraire approach is to elicit a value (in terms of the numeraire) of different levels of the commodity. We can

[10] Chapter 17 of Keynes (1936) discusses the issues here in the context of a whole economy and the determinant of interest rates. So important is this chapter that we were tempted to quote it in its entirety.

ask what amount of the commodity now is equally preferred to a fixed quantity of that commodity after a specified delay, translate these commodity amounts into the numeraire and calculate the implied rate of discount. Doing this for a range of commodities permits us to *compare* them in terms of the amount to which they are discounted.

A variant on the numeraire approaches is the Delayed Compensation Mechanism (DCM) introduced in Cubitt, McDonald and Read (in press) (Option 3). This method asks what amount of money, received after a delay, would be enough to make a person indifferent between their most preferred out of two options and their least preferred plus the money amount. The benefit of this method is that the options can be any goods available at any dates. The DCM approach does not require the outcomes to be divisible, nor do they have to satisfy nonsatiation. We can, for instance, ask how much money you would require to delay the receipt of a Ferrari by one year. So this approach opens up the possibility for a wide range of research questions in intertemporal choice that have previously been closed. However, the method does not directly give an estimate of the discount rate, instead giving comparative information about how heavily things are discounted.

A final option is to avoid the elicitation of discount rates altogether, but instead to see how the relative preference for two options changes over time

(Option 4). Take two goods that are equally preferred. To make it concrete, imagine someone is indifferent between two pens and three candies, if they could have them right now. If candies and pens are discounted at the same rate, then delaying the receipt of *both* the candies and the pen by one year will not change the relative preference between them: the person will still be indifferent. Now suppose pens are actually discounted by more than candies. In that case, if both the pens and the candies are delayed by one year, then the person will now prefer the candies (from the perspective of the present).

5.7 Prior Research into Commodity-Specific Discounting

We have just explained several potential pitfalls to avoid when investigating intertemporal choices for non-monetary goods. We have made clear our reservations about Option 1, which was to find out what proportion of its value something loses as a function of delay by matching quantities of an outcome over time. Nonetheless, the majority of existing studies have drawn on versions of Option 1. We will outline some of the major contributions to the literature and discuss whether they can be interpreted as revealing respondents' underlying discount rates.

Ubfal (2016) used Option 1 to elicit the discount rate for different goods, including money. The goods were chosen to be divisible in some way, such as meat, sugar and bottles of perfume, or else were described in monetary terms such as vouchers for school supplies and phone airtime. In other conditions, the goods were described in '$-worth' terms, as in '$1 worth of meat'. Ubfal found rates ranging from 66 per cent to 110 per cent per month, with discounting for money on the higher end of that range (90 per cent per month). The use of $-worth values aims to avoid problems with unobservable value functions, but does not fully accomplish this because the relevant money amount is actually the individual's willingness to pay for options, rather than the market price of those options[11]. For instance, if a person is given a choice between a Bic pen now (costing $1) and a Mont Blanc pen in one year (costing $500) and they take the Bic pen, this does not mean that they have a discount rate of 50,000 per cent for pens. More likely, they do not perceive the fancy pen to be 500 times as valuable to them as the Bic. Ubfal (2016) did attempt to control for this by eliciting a utility curvature estimate for money and then applying this across all goods, but of course the problem is not the utility function for money, but for the goods

[11] Though the market price will provide an indication of the population's willingness to pay, it cannot be assumed to accurately capture the individual willingness to pay for all respondents.

being discounted. Ubfal suggested an interesting next step would be to do this on a good-specific basis, and it still is.

Weatherly and Ferraro (2011) also used Option 1 to elicit and compare discount rates for different commodities. They explored outcomes such as debts owed and money won, federal legislation on education, having your ideal body and even finding the love of your life. They used the Area Under the Curve (AUC) method (see Myerson, Green and Warusawitharana, 2001), finding differences in discount rates across commodities. They made no concession for how much the outcomes are liked in the first place. However, their results indicate different domains of outcomes, where discounting is related within the domain but not between domains. This idea fits well with the suggestion that discounting will depend on a range of factors, such as visceral influences, satiation and so on, which might be expected to apply similarly within outcome domains.

Knapman and Tronde (2012) investigated closely related, yet not identical rewards. They found that different rewards from hypothetical airline schemes are discounted differently. However, their study demonstrates the difficulty in finding convincing divisibility in nonmonetary rewards, since in their quest to find divisible outcomes they have to assume constant marginal utility of the distance of a reward flight. For example, they calculate discounting on the basis of choices such as 'I would accept a free flight of __ km in 1 month in lieu of a free flight of 17000 kilometres in two months'. It is likely, for instance, that a free flight to the middle of the Pacific Ocean is worth less than half that to an airport even closer, and someone whose relatives live 500 kilometres away might not want a 2,000 kilometre flight.

Other scholars have taken an approach much more aligned with our Option 4 and investigated time preferences without eliciting a discount rate per se. Some examples include Read, Loewenstein and Kalyanaraman (1999), who examined preference reversals for high- and low-brow movie rentals and lotteries, and Read and van Leeuwen (1998), who investigated preference reversals for snack foods under hot and cold states. Another is Kirby and Herrnstein (1995), who used consumer goods and titrated the delay until preferences were reversed, providing evidence that different consumer goods have different discount rates – in Kirby and Herrnstein's case, either because the goods were differentially delayed in time, or because they were different goods. In none of these studies were the authors able to directly measure the discount rates for goods, since there was no common currency with which to do so. From their results, however, we can conclude that (for instance) junk food snacks have higher discount rates than healthy snacks (Read and Van Leeuwen, 1998), that low-brow movies have higher discount rates than high-brow ones (Read et al., 1999), and that immediate average-quality goods have lower discount rates than delayed higher-quality goods (Kirby and Herrnstein, 1995).

To get around the problems of yield and carrying costs, some innovative researchers have considered 'directly consumable rewards'. In the early work in this area, subjects were usually drug addicts and rewards were hypothetical. Bickel and Marsch (2001) reviewed the literature, and found that in most cases drug abusers discount the drug they abuse more than they discount money[12]. More recently, research has tended to focus on food choices with real incentives (i.e., the food or drink is consumed in the experiment). For example, Estle and colleagues (2007) compared the discounting of candy and drinks (i.e., soda and beer) and concluded that people discount the consumption of candy and drinks far more steeply than money. Reuben, Sapienza and Zingales (2010) compared the discounting of money and chocolate. To avoid satiation, they picked small, fancy chocolates that were not easily available on the market. They found the discount rates for money and chocolates are moderately correlated, but could not compare those rates directly. Tsukayama and Duckworth (2010) elicited discount rates for chips, candy, beer and money and measured impatience with

[12] This occurs when the street value of the drug is matched to the money value – a reasonable decision if the addicts will certainly use *at least that much* of the drug. Many studies involve addicts in treatment programs, who may actively wish not to use drugs in the future and so sharply devalue future drugs.

Area Under the Curve. Although they argued that directly consumable goods (i.e., chips, candy and beer) are discounted more steeply than money, they did not match utility across goods, so it is hard to know if the discount rates are truly comparable. However, Tsukayama and Duckworth (2010) drew a second line between directly consumable goods and money by showing that the correlation between discount rates of money and that of consumable goods are lower than the correlation of discount rates among different directly consumable goods. Food is a useful vehicle for studying intertemporal choice because the associated utility gain is received as soon as the food is consumed (and in the laboratory this can be tightly controlled). There are no concerns about storage or yield. There is another category of outcomes that fill this brief as well. These are *experienced* outcomes.

George Loewenstein has been highly innovative in this area, and has shown how apparently divergent discount rates for different goods can be consistent with the same underlying discount rate for utility. He has proposed the consumption of nonmonetary goods is associated with violations of additive separability, the assumption that the receipt of an outcome has no impact on the value placed on a future outcome. Loewenstein and Prelec (1993) proposed that people like their experiences to improve rather than to get worse over time, even if individually the experiences are the same. Later work (e.g., Frederick and

Loewenstein, 2008) showed that a preference for improvement is not universal, and identified a wide range of motives for sequence preferences: (1) anticipation and dread, (2) contrast effects, (3) extrapolation, (4) uncertainty, (5)opportunity cost, (6) pure time preference, (7) diminishing marginal utility, (8) equity among selves and (9) divide equally heuristic. Among them, 1 through 3 are factors favouring improving sequences, 4 through 6 are factors favouring deteriorating sequences and 7 through 9 are factors favouring flat sequences. Read and Powell (2002) similarly proposed (a demonstrated the existence of) a wide range of motives for sequence preferences, which also implied that a wide range of patterns could be observed.

Loewenstein (1987) and Rick and Loewenstein (2007) also proposed that future good and bad experiences could be augmented by anticipation. Knowing that you are going to have a positive experience in the future produces savouring, which may make you prefer to delay the experience rather than take it right away; and knowing you will have a negative experience can make you so anxious that you want to get it over with. Many researchers have since shown that dread, in particular, is a powerful motivator. People often want to get unpleasant experiences out of the way, and will even take a worse bad experience now in exchange for a better bad experience later (as predicted by Loewenstein, 1987). This produces high rates of apparent 'negative discounting'

(Casari and Dragone, 2011; Harris, 2012; Story et al., 2013). Loewenstein (1987) suggested these effects are a legitimate part of a behaviourally accurate discounting model.

> Since DU does not ordinarily incorporate savouring or dread, and since both of these factors attenuate devaluation, conventional estimates of discount rates should be biased downward, especially in situations where savouring or dread significantly affect devaluation … The bias in estimation of discount rates will be especially serious if savouring and dread are different for different categories of consumption. If this were the case, then the general assumption that 'the discount rate is independent of the category of consumption goods for which it is calculated' (Landsberger, 1971, p. 1351) would be invalid.

Clearly, scholars of intertemporal choice for consumer goods have struggled with the restrictions of divisibility and satiation. They have also been restricted by the fact that the consumer goods have market prices, and participants might have some ability to 'game' the choices in their experiments to take advantage of buying and selling opportunities in the market (much like we discussed early on with regard to the choice of SS and LL money amounts in the presence of a banking sector). One way to sidestep this latter concern is to use hypothetical outcomes, as in Loewenstein's (1987) paper, where he offered hypothetical kisses from one's favourite movie star. In fact, there are many domains where there exists no market for the outcomes, so that arbitrage is not possible. A good example is the domain of health.

In some respects, health outcomes are the best conceivable vehicle for finding the 'true' discount rate, if such a thing exists. They are directly experienced and cannot be bought or sold. Arguably, there is unlikely to be satiation in health. This is not to say that there is constant marginal utility from improvements to health: for example, improving the quality of life from a health state close to death may be much more valuable than improving an already good health state by the same amount. But it might be legitimate to assume linearity in the *duration* of a health state (holding its severity constant). For instance, imagine comparing a moderate headache lasting ten hours that starts now, and a moderate headache lasting twelve hours that starts in one week, where the intensity of the headache is identical. We might assume *temporal additivity*, which is that each hour of headache is just as bad as the other, and so we can obtain the total (dis)utility of the headache by multiplying its 'badness' by its duration. There are obviously limits to temporal additivity for health states,[13] but they are arguably less than those limits for any other domain.

Under this assumption, many researchers conduct matching studies that are, essentially, variants of Option 1, in which the duration of a health state is traded off against the delay until it would be experienced (Chapman and Elstein, 1995; Chapman, 1996; Van der Pol and Cairns, 2001). Unusually, Bleichrodt and Johanneson (2001) control for the nonlinearity of the duration of health states, to obtain discount rates that are not confounded by these nonlinearities. Overall, these papers show that many of the phenomena reported in intertemporal choices for money also are found for health states, including decreasing discounting with delay (cf. Bleichrodt et al., 2016, though McDonald et al., in press, found that the declining discount rates observed in their data were better explained by subadditive than by hyperbolic discounting).

The theoretical benefits of health as an outcome in the elicitation of intertemporal choice, in combination with its successful implementation in the studies we mentioned, make it an attractive option for the investigation of intertemporal choice. Health may allow us to learn about an individual's 'true' discount rate. However, this depends on whether there exists a 'true' discount

[13] For instance, each day of a lengthy illness, even if not associated with worsening symptoms, can nonetheless be worse than the day before.

rate to be measured. Evidence is mixed regarding the domain-generality of discount rates for health and for other outcomes. Evidence in favour was found by some (Moore and Viscusi, 1990; Cropper, Aydede and Portney, 1994) but not by others. Evidence that money is more discounted than health was reported by Cairns (1992), while a gain/loss effect was found by Cairns (1994) and by Hardisty and Weber (2009), with health gains discounted more than health losses. Papers directly estimating the correlation between health and money discounting typically find low but positive correlation (Chapman and Elstein, 1995; Chapman, 1996). However, for the reasons discussed earlier, since the health and financial discounting measures do not use the same numeraire, we cannot conclude absolutely that health and money are discounted at different rates.

There are also reasons to think that health behaviour does depend on a general discounting trait: Many studies show a correlation between time preference measured using money, and real-world health-related outcomes such as obesity, smoking and drug addiction (e.g., Chabris et al., 2008; Estle et al., 2007; Kirby, Petry and Bickel, 1999; Robles, 2010; Smith, Bogin and Bishai, 2005). The low correlations between health and monetary choices, however, underline a major issue in intertemporal choice research: these choices depend on a wide range of factors that researchers have only begun to investigate.

5.8 Conclusion

Researchers often assume the standard smaller sooner (SS) versus larger later (LL) paradigm – usually involving monetary amounts – provides a direct window on an underlying intertemporal discount rate for all goods, as if this is a 'trait' associated with the person, and not a rich interaction between the outcomes being evaluated, the way those outcomes (and delays separating outcomes) are represented, the situation in which the choice is made, as well as the person making the evaluation. We argued that a single experimental paradigm, or even a narrow class of such paradigms, may be insufficient to capture this richness, and researchers will need to find different ways to address both theoretical and empirical questions.

We focussed on the problem of how to measure time preference in a broader range of settings. Because the notion of time preference concerns utility from consumption, rather than the ratio between receipts, the standard methods do not help us very much. Even if we find, for instance, that someone is indifferent between ten bananas now and fifteen bananas in one month, we cannot say anything about their discount rate until we know how much consumption utility they will get from ten and fifteen bananas and how that utility will be distributed over time. We spent much of this chapter explaining the confusion and problems involved in making the link between raw receipts and consumption, as well as the problems in comparing rates of discounting elicited for different goods and in different contexts, and we suggested different approaches that researchers can take to alleviate these problems. As is clear from our examples, these methods need not be technically complex, but do require us to think of the measurement problem as being about something other than eliciting points of indifference between SS and LL. Indeed, as we discussed, relative discount rates can be assessed when two options are available at the same time, if we observe preferences between the options at different future points (and if preferences between them reverse). We also described a general method, the delayed compensation method (Cubitt et al., in press), for eliciting comparative rates of discounting. Other new methods include that of Attema et al. (2010, 2016), which enables discount rates to be measured directly without knowledge of the utility function.

While intertemporal choice has already received much attention, there is yet a great deal to be done. There is clearly a rich unexplored terrain for research in this area, and with clear reference to the normative theories, and insights from psychology and behavioural science, we will be able to exploit them and better understand a wide array of human decision-making.

chapter-references

5.9 References

Ainslie, G. (1975). Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control. *Psychological Bulletin*, 82(4), 463–496.

Ainslie, G. (2001). *Breakdown of Will*. Cambridge: Cambridge University Press.

Andersen, S., Harrison, G. W., Lau, M. I. and Rutström, E. E. (2008). Eliciting Risk and Time Preferences. *Econometrica,* 76(3), 583–618.

Andersen, S., Harrison, G. W., Lau, M. I. and Rutström, E. E. (2014). Discounting Behavior: A Reconsideration. *European Economic Review,* 71, 15–33.

Attema, A. E., Bleichrodt, H., Rohde, K. I. and Wakker, P. P. (2010). Time-Tradeoff Sequences for Analyzing Discounting and Time Inconsistency. *Management Science,* 56(11), 2015–2030.

Attema, A. E., Bleichrodt, H., Gao, Y., Huang, Z, and Wakker, P. P. (2016). Measuring Discounting Without Measuring Utility. *American Economic Review,* 106(6), 1476–1494.

Berns, G. S., Laibson, D. and Loewenstein, G. (2007). Intertemporal Choice – Toward an Integrative Framework. *Trends in Cognitive Sciences,* 11(11), 482–488.

Bickel, W. K. and Marsch, L. A. (2001). Toward a Behavioral Economic Understanding of Drug Dependence: Delay Discounting Processes. *Addiction*, 96(1), 73–86.

Bleichrodt, H., Gao Y. and Rohde, K. I. (2016). A Measurement of Decreasing Impatience for Health and Money. *Journal of Risk and Uncertainty*, 52(3), 213–231.

Bleichrodt, H. and Johannesson, M. (2001). Time Preference for Health: A Test of Stationarity Versus Decreasing Timing Aversion. *Journal of Mathematical Psychology*, 45(2), 265–282.

Casari, M. and Dragone, D. (2011). On Negative Time Preferences. *Economics Letters*, 111(1), 37–39.

Cairns, J. (1992). Discounting and Health Benefits: Another Perspective. *Health Economics*, 1(1), 76–79.

Cairns, J. A. (1994). Valuing Future Benefits. *Health Economics*, 3(4), 221–229.

Chabris, C. F., Laibson, D., Morris, C. L., Schuldt, J. P. and Taubinsky, D. (2008). Individual Laboratory-Measured Discount Rates Predict Field Behavior. *Journal of Risk and Uncertainty*, 37(2–3), 237–269.

Chapman, G. B. (1996). Temporal Discounting and Utility for Health and Money. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 771–791.

Chapman, G. B. and Elstein, A. S. (1995). Valuing the Future Temporal Discounting of Health and Money. *Medical Decision Making*, 15(4), 373–386.

Chung, S. and Herrnstein, R. J. (1961). Relative and Absolute Strengths of Response as a Function of Frequency of Reinforcement. *Journal of the Experimental Analysis of Animal Behavior*, 4(3), 267–272.

Cropper, M. L., Aydede, S. K. and Portney, P. R. (1994). Preferences for Life Saving Programs: How the Public Discounts Time and Age. *Journal of Risk and Uncertainty*, 8(3), 243–265.

Cubitt, R. P., McDonald, R. and Read, D. (in press). Time Matters Less When Outcomes Differ: Uni-modal Versus Cross-modal Comparisons in Intertemporal Choice. *Management Science*.

Cubitt, R. P. and Read, D. (2007). Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption? *Experimental Economics*, 10(4), 369–389.

Doyle, J. R. (2013). Survey of Time Preference, Delay Discounting Models. *Judgment and Decision Making*, 8(2), 116–135.

Ebert, J. E. and Prelec, D. (2007). The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future. *Management Science*, 53(9), 1423–1438.

Edgeworth, F. Y. (1879). The Hedonical Calculus. *Mind*, 4, 394–408.

Elster, J. (1979). *Ulysses and the Sirens: Studies in Rationaltiy and Irrationality.* Cambridge: Cambridge University Press.

Estle, S. J., Green, L., Myerson, J. and Holt, D. D. (2007). Discounting of Monetary and Directly Consumable Rewards. *Psychological Science*, 18(1), 58–63.

Faralla, V., Benuzzi, F., Lui, F., Baraldi, P., Dimitri, N. and Nichelli, P. (2015). Neural Correlates in Intertemporal Choice of Gains and Losses. *Journal of Neuroscience, Psychology, and Economics*, 8(1), 27–47

Fisher, I. (1930). *The Theory of Interest*. New York: Macmillan.

Frederick, S. and Loewenstein, G. (2008). Conflicting Motives in Evaluations of Sequences. *Journal of Risk and Uncertainty*, 37(2–3), 221–235.

Frederick, S., Loewenstein, G. and O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2), 351–401.

Gigliotti, G. and Sopher, B. (1997). Violations of Present-Value Maximization in Income Choice. *Theory and Decision*, 43, 45–69.

Glimcher, P. W., Kable, J. and Louie, K. (2007). Neuroeconomic Studies of Impulsivity: Now or Just as Soon as Possible? *American Economic Review*, 97(2), 142–147.

Green, L., Myerson, J. and Macaux, E. W. (2005). Temporal Discounting When the Choice Is Between Two Delayed Rewards. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(5), 1121–1133.

Halevy, Y. (2015). Time Consistency: Stationarity and Time Invariance. *Econometrica*, 83(1), 335–352.

Hardisty, D. J., Appelt, K. C. andWeber, E. U. (2013). Good or Bad, We Want It Now: Fixed-Cost Present Bias for Gains and Losses Explains Magnitude Asymmetries in Intertemporal Choice. *Journal of Behavioral Decision Making*, 26(4), 348–361.

Hardisty, D. J. and Weber, E. U. (2009). Discounting Future Green: Money Versus the Environment. *Journal of Experimental Psychology: General*, 138(3), 329–340.

Hariri, A. R., Brown, S. M., Williamson, D. E., Flory, J. D., de Wit, H. and Manuck, S. B. (2006). Preference for Immediate over Delayed Rewards Is Associated with Magnitude of Ventral Striatal Activity. *Journal of Neuroscience*, 26(51), 13213–13217.

Harris, C. R. (2012). Feelings of Dread and Intertemporal Choice. *Journal of Behavioral Decision Making*, 25(1), 13–28.

Holcomb, J. H. and Nelson, P. S. (1992). Another Experimental Look at Individual Time Preference. *Rationality and Society*, 4(2), 199–220.

Homer (Translated 1932). *The Odyssey*, T. E. Lawrence (translator). London: Macmillan Collectors Library (this edition published 2004).

Kable, J. W. (2013). Valuation, Intertemporal Choice, and Self-Control. In P. W. Glimcher and F. Ernst (Eds.), *Neuroeconomics: Decision Making and the Brain* (pp. 173–192): London: Academic Press.

Kable, J. W. and Glimcher, P. W. (2007). The Neural Correlates of Subjective Value During Intertemporal Choice. *Nature Neuroscience*, 10(12), 1625–1633.

Kable, J. W. and Glimcher, P. W. (2010). An 'as Soon as Possible' Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms. *Journal of Neurophysiology*, 103(5), 2513–2531.

Kahneman, D., Wakker, P. P. and Sarin, R. (1997). Back to Bentham? Explorations of Experienced Utility. *Quarterly Journal of Economics*, 112(2), 375–405.

Kayser, A. S., Allen, D. C., Navarro-Cebrian, A., Mitchell, J. M. and Fields, H. L. (2012). Dopamine, Corticostriatal Connectivity, and Intertemporal Choice. *Journal of Neuroscience*, 32(27), 9402–9409.

Keynes, J. M (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan (reprinted 2007).

Kinari, Y., Ohtake, F. and Tsutsui, Y. (2009). Time Discounting: Declining Impatience and Interval Effect. *Journal of Risk and Uncertainty*, 39(1), 87–112.

Kirby, K. N. and Herrnstein, R. J. (1995). Preference Reversals Due to Myopic Discounting of Delayed Reward. *Psychological Science*, 6(2), 83–89.

Kirby, K. N., Petry, N. M. and Bickel, W. K. (1999). Heroin Addicts Have Higher Discount Rates for Delayed Rewards Than Non-Drug-Using Controls. *Journal of Experimental Psychology: General*, 128(1), 78–87.

Knapman, A. and Tronde, J. (2011). *Temporal Discounting of Different Commodities in Airline Customer Reward Programs*. (MBA), Blekinge Institute of Technology. Retrieved from www.diva-portal.org/smash/get/diva2:829832/FULLTEXT01.pdf

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(2), 443–477.

Landsberger, M. (1971). Consumer Discount Rate and the Horizon: New Evidence. *Journal of Political Economy*, 79(6), 1346–1359.

Laury, S. K., McInnes, M. M. and Swarthout, J. T. (2012). Avoiding the Curves: Direct Elicitation of Time Preferences. *Journal of Risk and Uncertainty*, 44(3), 181–217.

LeBoeuf, R. A. (2006). Discount Rates for Time Versus Dates: The Sensitivity of Discounting to Time-Interval Description. *Journal of Marketing Research*, 43(1), 59–72.

Lempert, K. M. and Phelps, E. A. (2016). The Malleability of Intertemporal Choice. *Trends in Cognitive Sciences*, 20(1), 64–74.

Loewenstein, G. (1987). Anticipation and the Valuation of Delayed Consumption. *Economic Journal*, 97(387), 666–684.

Loewenstein, G. (1988). Frames of Mind in Intertemporal Choice. *Management Science*, 34(2), 200–214.

Loewenstein, G. and Prelec, D. (1992). Anomalies in Intertemporal Choice: Evidence and an Interpretation. *Quarterly Journal of Economics*, 107(2), 573–597.

Loewenstein, G. and Prelec, D. (1993). Preferences for Sequences of Outcomes. *Psychological Review*, 100(1), 91–108.

Loewenstein, G., Rick, S. and Cohen, J. D. (2008). Neuroeconomics. *Annual Review of Psychology*, 59, 647–672.

Loewenstein, G. and Sicherman, N. (1991). Do Workers Prefer Increasing Wage Profiles? *Journal of Labor Economics*, 9(1) 67–84.

Loewenstein, G. and Thaler, R. H. (1989). Anomalies: Intertemporal Choice. *Journal of Economic Perspectives*, 3(4), 181–193.

Magen, E., Dweck, C. S. and Gross, J. J. (2008). The Hidden-Zero Effect Representing a Single Choice as an Extended Sequence Reduces Impulsive Choice. *Psychological Science*, 19(7), 648–649.

Manuck, S. B., Flory, J. D., Muldoon, M. F. and Ferrell, R. E. (2003). A Neurobiology of Intertemporal Choice. In G. Loewenstein, D., Read and R. Baumeister

(Eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (pp. 139–172). New York: Russell Sage.

Manzini, P., Mariotti, M. and Mittone, L. (2010). Choosing Monetary Sequences: Theory and Experimental Evidence. *Theory and Decision*, 69(3), 327–354.

Mazur, J. E. (1987). An Adjusting Procedure for Studying Delayed Reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin and H. Rachlin (Eds.), *The Effect of Delay and of Intervening Events on Reinforcement Value: Quantitative Analyses of Behavior* (pp. 55–73). Hillsdale, NJ: Lawrence Erlbaum.

McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G. and Cohen, J. D. (2007). Time Discounting for Primary Rewards. *Journal of Neuroscience*, 27(21), 5796–5804.

McClure, S. M., Laibson, D. I., Loewenstein, G. and Cohen, J. D. (2004). Separate Neural Systems Value Immediate and Delayed Monetary Rewards. *Science*, 306(5695), 503–507.

McAlvanah, P. (2010). Subadditivity, Patience, and Utility: The Effects of Dividing Time Intervals. *Journal of Economic Behavior & Organization*, 76(2), 325–337.

McDonald, R., Chilton, S., Jones-Lee, M. W. and Metcalf, H. (in press). Evidence of Variable Discount Rates and Non-standard Discounting in Mortality Risk Valuation. *Journal of Environmental Economics and Management*

Moore, M. J. and Viscusi, W. K. (1990). Models for Estimating Discount Rates for Long-Term Health Risks Using Labor Market Data. *Journal of Risk and Uncertainty*, 3(4), 381–401.

Myerson, J., Green, L. and Warusawitharana, M. (2001). Area Under the Curve as a Measure of Discounting. *Journal of the Experimental Analysis of Behavior*, 76(2), 235–243.

O'Donoghue, T. and Rabin, M. (1999). Doing It Now or Later. *American Economic Review*, 89(1), 103–124.

O'Donoghue, T. and Rabin, M. (2000). The Economics of Immediate Gratification. *Journal of Behavioral Decision Making*, 13(2): 233–250

Phelps, E. S. and Pollak, R. A. (1968). On Second-Best National Saving and Game-Equilibrium Growth. *Review of Economic Studies*, 35(2), 185–199.

Plato. (trans. 1956). *Protagoras*. in B. Jowett and M. Ostwald (Trans.), G. Vlastos, (Ed.). Indianapolis: Bobbs-Merrill.

Read, D. (2001). Is Time-Discounting Hyperbolic or Subadditive? *Journal of Risk and Uncertainty*, 23(1), 5–32.

Read, D. (2004). Intertemporal Choice. In D. J. Koehler and N. Harvey (Eds.), *Blackwell Handbook of Judgment and Decision Making* (pp. 424–443). New Jersey: Blackwell.

Read, D. (2007). Experienced Utility: Utility Theory from Jeremy Bentham to Daniel Kahneman. *Thinking & Reasoning*, 13(1), 45–61.

Read, D., Frederick, S., Orsel, B. and Rahman, J. (2005). Four Score and Seven Years from Now: The Date/Delay Effect in Temporal Discounting. *Management Science*, 51(9), 1326–1335.

Read, D., Frederick, S. and Scholten, M. (2013). DRIFT: An Analysis of Outcome Framing in Intertemporal Choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(2), 573–588.

Read, D., Loewenstein, G. and Kalyanaraman, S. (1999). Mixing Virtue and Vice: Combining the Immediacy Effect and the Diversification Heuristic. *Journal of Behavioral Decision Making*, 12(4), 257–273.

Read, D., Olivola, C. Y. and Hardisty, D. (2016). The Value of Nothing: Asymmetric Attention to Opportunity Costs Drives Intertemporal Decision Making, *Management Science*, dx.doi.org/10.1287/mnsc.2016.2547

Read, D. and Powell, M. (2002). Reasons for Sequence Preferences. *Journal of Behavioral Decision Making*, 15(5), 433–460.

Read, D. and Read, N. L. (2004). Time Discounting over the Lifespan. *Organizational Behavior and Human Decision Processes*, 94(1), 22–32.

Read, D. and Roelofsma, P. H. (2003). Subadditive Versus Hyperbolic Discounting: A Comparison of Choice and Matching. *Organizational Behavior and Human Decision Processes*, 91(2), 140–153.

Read, D. and Van Leeuwen, B. (1998). Predicting Hunger: The Effects of Appetite and Delay on Choice. *Organizational Behavior and Human Decision Processes*, 76(2), 189–205.

Reuben, E., Sapienza, P. and Zingales, L. (2010). Time Discounting for Primary and Monetary Rewards. *Economics Letters*, 106(2), 125–127.

Rick, S. and Loewenstein, G. (2008). The Role of Emotion in Economic Behavior. In M. Lewis, J. M. Haviland-Jones and L. F. Barrett (Eds.), *Handbook of Emotions* (pp. 138–156). New York: Guilford Press.

Robles, E. (2010). Delay and Loss of Subjective Value. *Journal of Behavior, Health & Social Issues*, 2(1), 105–118.

Roelofsma, P. H. (1996). Modelling Intertemporal Choices: An Anomaly Approach. *Acta Psychologica,* 93(1), 5–22.

Rubinstein, A. (2003). 'Economics and Psychology'? The Case of Hyperbolic Discounting. *International Economic Review*, 44(4), 1207–1216.

Samuelson, P. A. (1937). A Note on Measurement of Utility. *Review of Economic Studies,* 4(2), 155–161.

Sayman, S. and Öncüler, A. (2009). An Investigation of Time Inconsistency. *Management Science,* 55(3), 470–482.

Scholten, M. and Read, D. (2006). Discounting by Intervals: A Generalized Model of Intertemporal Choice. *Management Science*, 52(9), 1424–1436.

Scholten, M. and Read, D. (2010). The Psychology of Intertemporal Tradeoffs. *Psychological Review,* 117(3), 925–944.

Scholten, M. and Read, D. (2013). Time and Outcome Framing in Intertemporal Tradeoffs. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 39(4), 1192–1212.

Scholten, M., Read, D. and Sanborn, A. (2016). Cumulative Weighing of Time in Intertemporal Tradeoffs. *Journal of Experimental Psychology: General,* 145(9), 1177–1205.

Sellitto, M., Ciaramelli, E. and di Pellegrino, G. (2010). Myopic Discounting of Future Rewards After Medial Orbitofrontal Damage in Humans. *Journal of Neuroscience*, 30(49), 16429–16436.

Shapiro, J. M. (2005). Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle. *Journal of Public Economics*, 89(2), 303–325.

Smith, C. T., Wallace, D. L., Dang, L. C., Aarts, E., Jagust, W. J., D'Esposito, M. and Boettiger, C. A. (2016). Modulation of Impulsivity and Reward Sensitivity in Intertemporal Choice by Striatal and Midbrain Dopamine Synthesis in Healthy Adults. *Journal of Neurophysiology*, 115(3), 1146–1156.

Smith, P. K., Bogin, B. and Bishai, D. (2005). Are Time Preference and Body Mass Index Associated? Evidence from the National Longitudinal Survey of Youth. *Economics & Human Biology*, 3(2), 259–270.

Sopher, B. and Sheth, A. (2006). A Deeper Look at Hyperbolic Discounting. *Theory and Decision*, 60(2–3), 219–255.

Story, G. W., Vlaev, I., Seymour, B., Winston, J. S., Darzi, A. and Dolan, R. J. (2013). Dread and the Disvalue of Future Pain. *PLoS Computational Biology*, 9(11), e1003335.

Strotz, R. H. (1955). Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies*, 23(3), 165–180.

Takahashi, T. (2005). Loss of Self-Control in Intertemporal Choice May Be Attributable to Logarithmic Time-Perception. *Medical Hypotheses*, 65(4), 691–693.

Thaler, R. (1981). Some Empirical Evidence on Dynamic Inconsistency. *Economics Letters*, 8(3), 201–207

Tsukayama, E. and Duckworth, A. L. (2010). Domain-Specific Temporal Discounting and Temptation. *Judgment and Decision Making*, 5(2), 72–82.

Tversky, A., Sattath, S. and Slovic, P. (1988). Contingent Weighting in Judgment and Choice. *Psychological Review*, 95(3), 371–384.

Ubfal, D. (2016). How General Are Time Preferences? Eliciting Good-Specific Discount Rates. *Journal of Development Economics*, 118, 150–170.

Urminsky, O. and Zauberman, G. (2015). The Psychology of Intertemporal Preferences. In G. Wu and G. Keren (Eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making* (pp. 141–181). Malden, MA: Wiley-Blackwell

Van der Pol, M. and Cairns, J. (2001). Estimating Time Preferences for Health Using Discrete Choice Experiments. *Social Science & Medicine*, 52(9), 1459–1470.

Weatherly, J. N. and Ferraro, F. R. (2011). Executive Functioning and Delay Discounting of Four Different Outcomes in University Students. *Personality and Individual Differences*, 51(2), 183–187.

Weber, E. U., Johnson, E. J., Milch, K. F., Chang, H., Brodscholl, J. C. and Goldstein, D. G. (2007). Asymmetric Discounting in Intertemporal Choice a Query-Theory Account. *Psychological Science*, 18(6), 516–523.

Wu, C. Y. and He, G.-B. (2012). The Effects of Time Perspective and Salience of Possible Monetary Losses on Intertemporal Choice. *Social Behavior and Personality*, 40, 1645–1653.

Yates, J. F. and Watts, R. A. (1975). Preferences for Deferred Losses. *Organizational Behavior and Human Performance*, 13(2), 294–306.

Zauberman, G., Kim, B. K., Malkoc, S. A. and Bettman, J. R. (2009). Discounting Time and Time Discounting: Subjective Time Perception and Intertemporal Preferences. *Journal of Marketing Research*, 46(4), 543–556.