Life Course Centre

WORKING PAPER SERIES

No. 2024-05

March 2024

Time Preference

Stephen L. Cheung

The Australian Research Council Centre of Excellence for Children and Families over the Life Course
Phone +61 7 3346 7477 **Email** lcc@uq.edu.au
lifecoursecentre.org.au

Australian Government
Australian Research Council

The University of Queensland Australia

The University of Sydney

The University of Western Australia

The University of Melbourne

Research Summary

Why was the research done?

The study of time preference is concerned with decisions involving benefits and costs at different points in time. Many of the most consequential choices over the life course have this feature. This article was written to provide a high-level, non-technical overview of both conceptual and practical issues that arise in the measurement of time preference in experiments and surveys.

What were the key findings?

While the basic structure of a time preference experiment is simple, there are many variables that must be specified by the researcher, and these will depend on the purpose of the measurement. They include the model of time preference applicable to a situation, the use of incentives, the choice of elicitation procedures and delay structures, and whether and how to control for non-linear valuation of outcomes. Where full-scale measurement is not feasible or desirable, a variety of qualitative survey-based measures have also been validated.

What does this mean for policy and practice?

In making design choices for a time preference study, researchers should consider among other things: whether time preference is a focal behaviour, key explanatory variable, or other covariate for the study at hand, whether the applicable time preference construct entails present bias and self-control problems, and the intended disciplinary audience. There is currently no methodological gold standard for the measurement of time preferences.

Citation

Cheung, S.L. (2024). 'Time preference', Life Course Centre Working Paper Series, 2024-05. Institute for Social Science Research, The University of Queensland.

The authors

Stephen L. Cheung
The University of Sydney, School of Economics
Email: Stephen.Cheung@sydney.edu.au

Acknowledgements/Funding Sources

This article was prepared as an entry for the *Encyclopedia of Experimental Social Science*.

Time preference*

Stephen L. Cheung†

March 2024

The study of *time preference* – also known in economics as *intertemporal choice*, and in psychology as *delay discounting* – is concerned with decisions that involve benefits and costs at different points in time. Applications range from the prosaic (watching football on TV instead of working on my paper) to the most consequential choices an individual makes over the life course (education, healthy lifestyle, and saving for retirement). This article provides a high-level overview of conceptual and practical issues that arise in the experimental measurement of time preference. The exposition is non-technical but abstract, and deliberately sets aside issues of statistical estimation. Suggestions for further reading are provided.

Models of discounting

A basic stylised fact of time preference is that most people (and animals) are *impatient*: they would prefer a sooner reward to the same-sized reward at a later date. Stated differently, the subjective value of a reward is *discounted* as it recedes into the future. This discounting may take different forms, and it is helpful to be clear from the outset on some distinctions between them. Figure 1 depicts a decision-maker faced with the choice between a *smaller-sooner* reward (x at date t) or a *larger-later* alternative ($y > x$ at date $t + \tau$). Of particular interest for expositional purposes are the cases where she faces this choice at date 0, and where she faces it at date t (in which case the sooner option is available immediately).

One possibility is that every period of delay is discounted at a constant rate. In traditional economics, this is the model of *exponential discounting* (Samuelson, 1937). It has the important implication that a chooser's preferences do not change merely as a result of the passing of time (*time consistency*). In the top-left panel of Figure 1 the chooser is more impatient, consistently preferring smaller-sooner at all dates between 0 and t (the dotted line is everywhere above the dashed one); in the top-right she is less impatient, consistently preferring larger-later.

*Entry prepared for the *Encyclopedia of Experimental Social Science*.

†The University of Sydney, School of Economics, Social Sciences Building A02, Sydney NSW 2006, Australia (Stephen.Cheung@sydney.edu.au).

Figure 1: Preferences under alternative discounting models

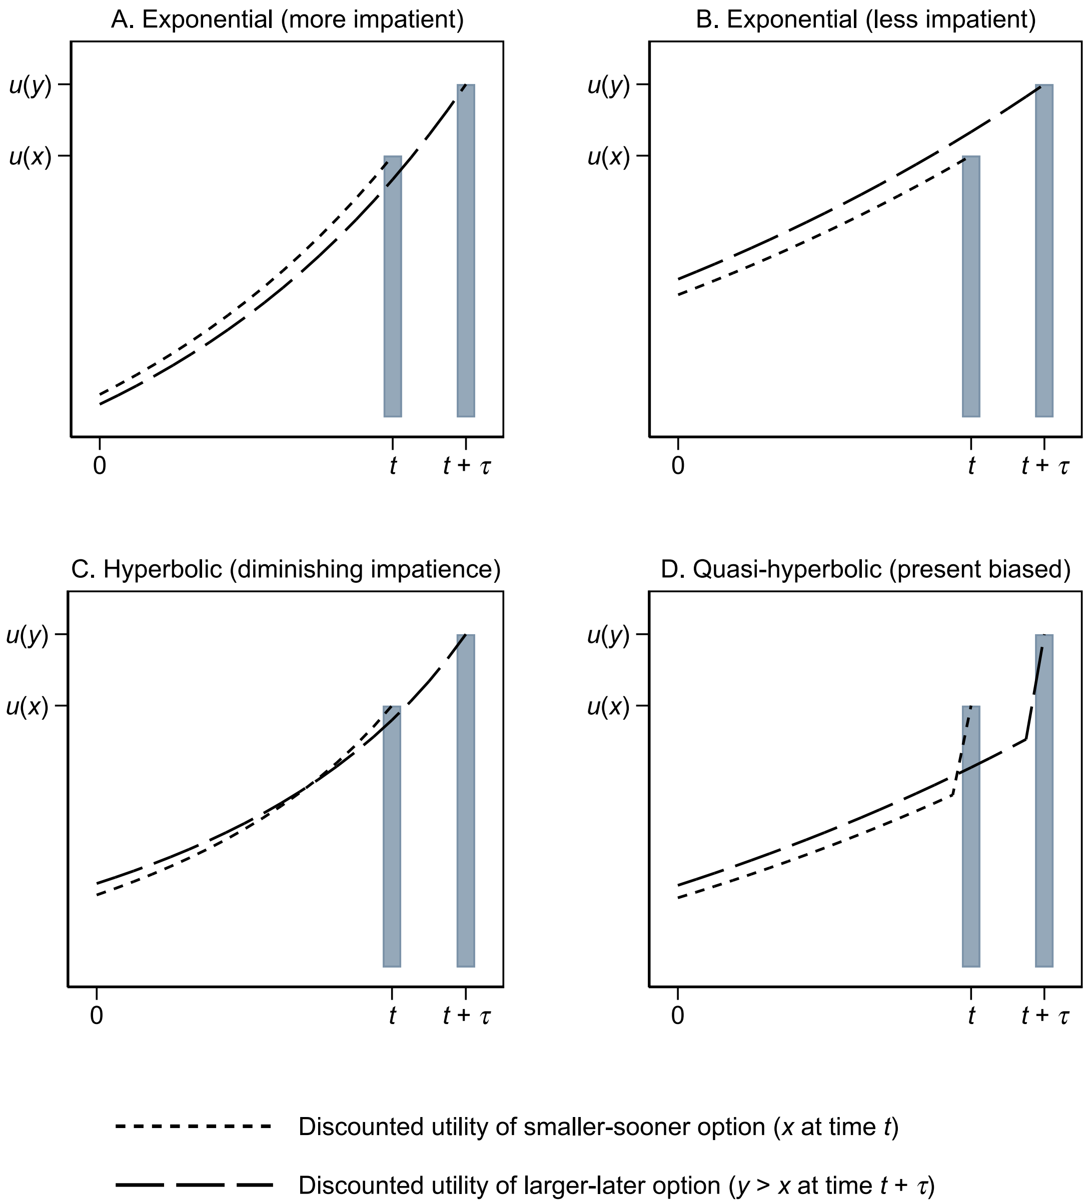

Notes: Preferences between smaller-sooner and larger-later rewards under alternative models of discounting. The height of each bar represents the undiscounted utility (subjective value) of an immediately-available reward, the horizontal axes denote chronological time, and the curves represent discounted utilities of future rewards as viewed from different points in time. (A, B) An exponential chooser discounts every period of delay at a constant rate, and thus has a consistent preference for either the sooner or later reward. (C) A hyperbolic chooser discounts more steeply as a reward becomes imminent, and may thus switch from preferring the later to the sooner reward as its date approaches. (D) A quasi-hyperbolic chooser discounts exponentially except in the presence of an immediate reward, and would thus only switch once a sooner reward is immediately available.

Any other form of discounting may yield *time-inconsistent* behaviour (Strotz, 1955): a chooser's preference might change over time, even without any change in information or circumstances. In psychology, the leading alternative to the exponential model is *hyperbolic discounting* (Mazur, 1987; Myerson and Green, 1995). This is characterised by *diminishing impatience*: the chooser is more impatient for the near future than for the far future. In the bottom-left of Figure 1 the curves become steeper as the respective rewards are more temporally proximate. As a result, the chooser initially prefers larger-later when both alternatives are far in the future, but at some point (prior to t) she exhibits a *preference reversal*, switching to smaller-sooner.

In behavioural economics, the leading model is *quasi-hyperbolic discounting* (Laibson, 1997). This is characterised by *present bias*: a chooser is more impatient specifically for the present than for all future periods. In the bottom-right of Figure 1, both curves are exponential in shape except at the moment of reward: each period is discounted at a constant rate, but there is a premium attached to an immediate reward.[1] This chooser prefers larger-later at all dates prior to t, but switches to smaller-sooner at t once presented with an immediate reward.

A key take-away from this discussion is that individual differences in patience – of the type described by exponential discounting – predict that people make more near- or far-sighted choices in general, but do not specifically predict that they are impulsive, succumb to temptation, or have self-control problems. Such predictions emerge only from a model that generates time-inconsistent preferences. Thus if a researcher wishes to predict or explain such behaviours, it is essential to have an experimental design that allows for such preferences to be expressed.

Measurement of time preference

The basic structure of a time preference experiment consists of a series of smaller-sooner vs. larger-later choices in which the amounts and dates are varied systematically. However, there are many variables that must be specified by the researcher, and these choices will depend on the purpose of the measurement. This section discusses some key issues.

Incentives

A fundamental tenet of experimental economics is that choices should be *incentive compatible*: it should be in participants' best interests to report their preferences truthfully. In practice, this usually means that one trial is randomly-selected for each participant, who receives the reward chosen in that trial on the date specified.[2] Rewards are typically amounts of money, but can also be real consumption or aversive "rewards" such as costly effort (Augenblick et al., 2015). By contrast in psychology, hypothetical choices and flat payments are generally accepted.

[1] Formally, all periods other than the present are subject to a uniform additional discount. The exponential model is the special case in which this additional discount is absent.

[2] In larger studies, it is acceptable for a random subset of participants (say one-in-ten) to be rewarded in this way.

In time preference experiments, the use of incentives introduces specific complications. First, both sooner and later rewards should be equally credible, free from risk, and have equal transaction costs of claiming rewards. If not, participants' choices will be biased – usually toward sooner choices from which the researcher will infer a high degree of impatience. These concerns are especially pressing when sooner rewards are potentially immediate: if immediate rewards are specifically more credible, less risky and more convenient to redeem, the extent of present bias would also be overstated. Second, real incentives are infeasible for certain reward types (e.g., health outcomes) and time horizons (on the scale of decades). Finally, incentives can either complicate or even invalidate some elicitation methods, favoured by psychologists, that are more parsimonious than pre-specified sets of binary choices.

Some economists have challenged the importance of incentives in time preference experiments. In recent meta-analyses, Matousek et al. (2022) find no effect of hypothetical rewards for estimates of impatience, while Cheung et al. (2023) find no effect for present bias. Brañas-Garza et al. (2023) find little difference between hypothetical, pay-one-in-ten, and real incentives in a series of laboratory, field, and online experiments. Nonetheless, until this debate is more settled, economists are advised to continue using real incentives where possible, especially where time preference is either the focal behaviour in a study or a key explanatory variable.

Elicitation procedures

Recall the framework of Figure 1: a participant chooses at date 0 between x at t, or $y > x$ at $t + \tau$. Suppose to begin that the researcher has committed to an exponential model, which assumes that every period of delay is discounted at a constant rate. In that case, the problem simplifies to identifying each participant's *discount rate*, thereby differentiating between cases such as those depicted in the two top panels of Figure 1.

A standard design in economics is the *choice list* (Coller and Williams, 1999; Harrison et al., 2002), an ordered list of binary choices in which the reward dates and one of the magnitudes are fixed, while the other magnitude incrementally varies. For example, the sooner reward x can be fixed while the value of y progressively improves. At an initial position where y and x are equal (so there is no reward for waiting), essentially all participants are expected to choose sooner. Then as y increases, the least impatient start choosing the later option while the most impatient hold out for larger values of y. In this way, it possible to find a point of *indifference* for each participant, where x and y are such that they are just willing to tolerate the delay; it is from this indifference value that their discount rate can be inferred.

It might seem cumbersome to elicit a point of indifference using a fixed set of binary choices. An attractive alternative could be to ask directly, for example (this time fixing y) "What value of x at t would make you equally happy as y at $t + \tau$?". This is known as a *matching* task. While straightforward to implement in hypothetical choices, matching is seldom used in incentivised time preference experiments (a notable exception is Benhabib et al., 2010) as the mechanism for its incentivisation (Becker et al., 1964) is complicated to explain to participants.[3]

Another approach to streamlining the elicitation procedure is to present items adaptively in response to choices in past trials. Thus if a participant chose smaller-sooner on the first trial the value of y can be increased in the next to make the larger-later option more attractive, while if she chose larger-later y can be decreased to make it less attractive. In this way it is possible to "hone in" to a point of indifference using only a handful of trials. This is known as a *staircase* or *titration* design. Unfortunately, some economists are sceptical of this approach as they are wary of its incentive compatibility. In particular, it is conjectured that a participant might perceive an opportunity to manipulate the mechanism by misreporting their preference on earlier trials in order to influence the choices offered in later ones. Once again, this illustrates the challenges and trade-offs that result from the use of incentives.

Delay structures

Moving beyond the study of exponential discounting, it is necessary to introduce variation in t and τ, the *front-end delay* from the date of the experiment to the sooner reward date, and the *back-end delay* between sooner and later rewards. In particular, variation in t (specifically whether or not a sooner reward is available immediately) is required to identify present bias, while variation in τ is needed to identify diminishing impatience in a hyperbolic model.

Researchers in psychology focus on the hyperbolic model and its generalisations, and do not insist on incentives. They typically study hypothetical choices, set $t = 0$ (such that sooner outcomes are immediate), and utilise extensive variation in τ (sometimes up to 25 years). By contrast, economists tend to focus on the quasi-hyperbolic model, and are adamant of the need for incentives. Their studies often compare choice situations with $t = 0$ to ones with $t > 0$, going to great pains to equalise credibility, risk, and transaction costs. Since the quasi-hyperbolic model implies constant discounting away from the present, these designs do not require extensive variation in τ, while the practicalities of real incentives limit τ to at most one year. Because of these contrasting disciplinary norms, there is a surprising paucity of studies that allow for a comparison of fit between economists' and psychologists' preferred models.

Longitudinal designs

In experiments to identify present bias, participants are typically presented with choices both with and without front-end delay in the course of a single session, for example:

- Choice 1 (asked at date 0): Would you prefer x at date t, or y at date $t + \tau$?
- Choice 2 (asked at date 0): Would you prefer x at date 0, or y at date τ?

[3] Indeed, some instructions for this mechanism amount to describing a fictitious choice list (Healy, 2020).

These choices differ only in that both options are shifted forward by t periods in Choice 2. Supposing a participant chooses larger-later in Choice 1 and smaller-sooner in Choice 2, it is *assumed* that at date t she would also prefer an immediate x at t over y at $t + \tau$, thus exhibiting a present-biased preference reversal. A more direct approach is to return at date t to ask:

- Choice 2A (asked at date t): Would you prefer x at date t, or y at date $t + \tau$?

Comparing Choice 1 to Choice 2A, the participant faces the exact same trade-off at two different points in time, a more direct test of present bias. This is known as a *longitudinal design* (Read et al., 2012; Halevy, 2015; Cheung et al., 2022). While compelling, such designs are logistically more complex and introduce a risk of attrition. As such, this approach is only recommended for studies where documenting the robustness of present bias is a key focus.

Utility curvature

Up to now, it has been assumed that the preference between smaller-sooner and larger-later is driven entirely by the attitude toward delay. However these alternatives clearly differ on two dimensions, namely timing (sooner-later) and magnitude (smaller-larger), each of which may influence the preference. Suppose for example that a chooser prefers one apple today to two apples tomorrow. This preference may derive not only from the fact that delay is aversive (time discounting), but also because a second apple provides less additional satisfaction than the first (diminishing marginal utility). If researchers overlook the influence of the latter (by assuming that utility is linear in reward magnitude) they may overstate the influence of the former, thereby concluding that the chooser is more impatient than she really is.

Correcting for non-linear utility is one of the most significant recent advances in economic experiments on time preference (Andersen et al., 2008), but has also generated considerable controversy (Andreoni and Sprenger, 2012). The issues are technical, and Cheung (2016) provides a review focused specifically on this topic. In broad terms there are two distinct approaches, and they do not reach the same conclusion. The first method "borrows" information on the utility function from a separate domain (namely choice under risk) and superimposes it upon choices in a discounting task (Andersen et al., 2008). Because utility in choice under risk is rather concave, this results in a substantial lowering of estimated discount rates. The second approach measures the utility function directly from choice over time, finding that it is quite close to linear (Andreoni and Sprenger, 2012; Abdellaoui et al., 2013; Cheung, 2020). The implication is that near-linear utility will have minimal effect upon estimated discount rates (whereas imposing a risk-based measure would result in an overcorrection).

Survey-based measures

In studies where it is not a focal behaviour or key explanatory variable, researchers may wish to include time preference as a covariate despite not having the necessary time or resources for a full-strength measurement. A number of hypothetical and qualitative survey measures have been validated against incentivised experiments for use in such situations.

Vischer et al. (2013) study the item "Are you generally an impatient person, or someone who always shows great patience?" from the German Socio-Economic Panel (SOEP). It significantly predicts impatience in an incentivised choice list, though the raw correlation is a modest 0.15. Falk et al. (2023) use a lab experiment to validate a battery of survey measures of risk, time, and social preferences. For discounting, the best predictor is simply a hypothetical-choice version of the experiment itself, although reducing this to a five-step staircase sacrifices minimal explanatory power. They also find that a qualitative self-assessment, "How willing are you to give up something that is beneficial for you today in order to benefit more from that in the future?", has additional explanatory power beyond the quantitative measure.

In both studies discussed above, the incentivised experiment had only one pair of reward dates, and thus could not distinguish between different forms of discounting. Pinger (2017) studies an item from the SOEP thought to capture present bias: "I live for today and do not think about tomorrow". This is found to predict incentivised experimental choices in which the sooner payment is immediate, but not ones in which it is deferred.

Conclusion

This article has highlighted a range of considerations that arise in the experimental measurement of time preference. In making design choices for their own studies, researchers should consider among other things: the role of time preference in the study (as focal behaviour, key explanatory variable, or other covariate), the relevant time preference construct (generic impatience, diminishing impatience, or present bias), and the intended disciplinary audience. As noted over fifteen years ago by Chabris et al. (2008), it remains the case that there is no methodological gold standard for the measurement of time preference.

Further reading

This article has deliberately set aside more technical issues of statistical estimation. Key references on this (for economists) include Andersen et al. (2008, 2014) and Andreoni and Sprenger (2012). Frederick et al. (2002) is a classic review article in economics. Cheung (2016) provides a concise account of recent methodological developments, with a focus on the utility function. Kable (2014) covers developments in neuroscience.

References

Abdellaoui, M., Bleichrodt, H., l'Haridon, O., Paraschiv, C., 2013. Is there one unifying concept of utility? An experimental comparison of utility under risk and utility over time. Management Science 59, 2153–2169.

Andersen, S., Harrison, G.W., Lau, M.I., Rutström, E.E., 2008. Eliciting risk and time preferences. Econometrica 76, 583–618.

Andersen, S., Harrison, G.W., Lau, M.I., Rutström, E.E., 2014. Discounting behavior: A reconsideration. European Economic Review 71, 15–33.

Andreoni, J., Sprenger, C., 2012. Estimating time preferences from convex budgets. American Economic Review 102, 3333–3356.

Augenblick, N., Niederle, M., Sprenger, C., 2015. Working over time: Dynamic inconsistency in real effort tasks. Quarterly Journal of Economics 130, 1067–1115.

Becker, G.M., DeGroot, M.H., Marschak, J., 1964. Measuring utility by a single-response sequential method. Behavioral Science 9, 226–232.

Benhabib, J., Bisin, A., Schotter, A., 2010. Present-bias, quasi-hyperbolic discounting, and fixed costs. Games and Economic Behavior 69, 205–223.

Brañas-Garza, P., Jorrat, D., Espín, A.M., Sánchez, A., 2023. Paid and hypothetical time preferences are the same: Lab, field and online evidence. Experimental Economics 26, 412–434.

Chabris, C.F., Laibson, D.I., Schuldt, J.P., 2008. Intertemporal choice, in: Durlauf, S.N., Blume, L.E. (Eds.), The New Palgrave Dictionary of Economics (Second Edition). Palgrave Macmillan.

Cheung, S.L., 2016. Recent developments in the experimental elicitation of time preference. Journal of Behavioral and Experimental Finance 11, 1–8.

Cheung, S.L., 2020. Eliciting utility curvature in time preference. Experimental Economics 23, 493–525.

Cheung, S.L., Tymula, A., Wang, X., 2022. Present bias for monetary and dietary rewards. Experimental Economics 25, 1202–1233.

Cheung, S.L., Tymula, A., Wang, X., 2023. A meta-analysis of quasi-hyperbolic discounting. University of Sydney.

Coller, M., Williams, M.B., 1999. Eliciting individual discount rates. Experimental Economics 2, 107–127.

Falk, A., Becker, A., Dohmen, T., Huffman, D., Sunde, U., 2023. The preference survey module: A validated instrument for measuring risk, time, and social preferences. Management Science 69, 1935–1950.

Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: A critical review. Journal of Economic Literature 40, 351–401.

Halevy, Y., 2015. Time consistency: Stationarity and time invariance. Econometrica 83, 335–352.

Harrison, G.W., Lau, M.I., Williams, M.B., 2002. Estimating individual discount rates in Denmark: A field experiment. American Economic Review 92, 1606–1617.

Healy, P.J., 2020. Explaining the BDM – or any random binary choice elicitation mechanism – to subjects. Ohio State University.

Kable, J.W., 2014. Valuation, intertemporal choice, and self-control, in: Glimcher, P.W., Fehr, E. (Eds.), Neuroeconomics (Second Edition). Academic Press, pp. 173–192.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112, 443–478.

Matousek, J., Havranek, T., Irsova, Z., 2022. Individual discount rates: A meta-analysis of experimental evidence. Experimental Economics 25, 318–358.

Mazur, J.E., 1987. An adjusting procedure for studying delayed reinforcement, in: Commons, M.L., Mazur, J.E., Nevin, J.A., Rachlin, H. (Eds.), The Effect of Delay and of Intervening Events on Reinforcement Value. Psychology Press, pp. 55–73.

Myerson, J., Green, L., 1995. Discounting of delayed rewards: Models of individual choice. Journal of the Experimental Analysis of Behavior 64, 263–276.

Pinger, P.R., 2017. Predicting experimental choice behavior and life outcomes from a survey measure of present bias. Economics Bulletin 37, 2162–2172.

Read, D., Frederick, S., Airoldi, M., 2012. Four days later in Cincinnati: Longitudinal tests of hyperbolic discounting. Acta Psychologica 140, 177–185.

Samuelson, P.A., 1937. A note on the measurement of utility. Review of Economic Studies 4, 155–161.

Strotz, R.H., 1955. Myopia and inconsistency in dynamic utility maximization. Review of Economic Studies 23, 165–180.

Vischer, T., Dohmen, T., Falk, A., Huffman, D., Schupp, J., Sunde, U., Wagner, G.G., 2013. Validating an ultra-short survey measure of patience. Economics Letters 120, 142–145.