

A meta-analysis of quasi-hyperbolic discounting*

Stephen L. Cheung
The University of Sydney

Agnieszka Tymula
The University of Sydney

Xueting Wang[†]
RMIT University

September 2025

Abstract

This paper reports a meta-analysis of 86 studies of the two parameters of quasi-hyperbolic (β - δ) discounting, the dominant model of self-control failures in behavioural economics. The central tenet of the model is that decision-makers have a “present-bias”, β , for immediate rewards, on top of standard exponential discounting of the future, δ . After correcting for selective reporting, we obtain a meta-analytic estimate of β for money of 0.938, with 95% confidence interval [0.905, 0.972]. For non-monetary rewards, our estimate of β is 0.750, with 95% confidence interval [0.643, 0.857], and there is no evidence of selective reporting. We find that present bias for real effort, while stronger than for money, is weaker than for other rewards such as real consumption.

Keywords: quasi-hyperbolic discounting; present bias; discount factor; beta-delta model; meta-analysis.

JEL codes: C91, D12, D80, D91.

Conflict of interest statement: The authors have no conflicts of interest to declare.

*We thank the Glimcher Lab at New York University and the Kable Lab at University of Pennsylvania for hosting Wang’s visit while working on this project and for helpful comments in the early stages of data analysis. We thank participants in the 2019 Sydney Workshop on Experimental Economics and Theory, 2021 Economic Science Association Annual Meeting, UNSW Behavioural Insights for Business and Policy Luncheon and the 2020 Society for Neuroeconomics Annual Meeting for valuable discussions. We are grateful to Alex Berger, Xinyue Cai, Matthew Elias, Zhenlin Kang, Alex Kiefer, Lexin Liang, Sura Majeed, Adam Teperski, and Olivia Ru for their excellent research assistance. Cheung acknowledges support from the Australian Research Council (DP160101794). Tymula acknowledges support from the University of Sydney bridging grant (G199763). Wang acknowledges PhD fellowship support from the Australian Research Council Centre of Excellence for Children and Families over the Life Course (CE140100027 and CE200100025).

[†]Corresponding author. School of Economics, Finance, and Marketing, RMIT University (xueting.wang@rmit.edu.au).

1 Introduction

Self-control is a crucial attribute in the realm of managerial economics, as it plays a pivotal role in decision-making processes and overall organisational success. Managers must exercise self-control when making myriad choices that have profound implications for their companies. The ability to resist immediate gratification in favour of long-term strategic objectives is essential in allocating resources efficiently and making optimal investment decisions. Moreover, self-control is vital in maintaining discipline within teams, ensuring adherence to established budgets and timelines, and fostering a culture of accountability. From a managerial perspective, the practice of self-control not only enhances individual leadership effectiveness but also cultivates an environment conducive to sustained growth and competitiveness.

Regrettably, decision-makers often fail to adhere to their plans, instead succumbing to immediate temptations at the expense of long-run benefits (DellaVigna and Malmendier, 2006; Kaur et al., 2015; Laibson et al., 1998). This is especially the case in decisions involving work and education (e.g., keeping to a schedule versus procrastinating), finances (e.g., saving and investing for the future versus consuming today for pleasure), and health (e.g., maintaining a healthy lifestyle and diet versus backsliding).

In behavioural economics and other related disciplines, such behaviours are explained in terms of *time-inconsistent preferences*. The dominant model of quasi-hyperbolic (or β - δ) discounting (Laibson, 1997) assumes that decision makers have a “present bias” for immediate gratification, such that when evaluating trade-offs that involve an immediate reward, the value of *all* delayed rewards is down-weighted by a factor $\beta < 1$ (on top of standard exponential discounting of the future). While this model is widely invoked to explain problematic behaviours across diverse domains, the extent to which it is supported by the available empirical evidence has been the subject of some controversy.

In the early 2000s, it was widely accepted as a stylised fact that people are present biased (Frederick et al., 2002) even though a precise estimate of β was not available. More general evidence of non-exponential discounting dates back at least to the early 1980s. Thaler (1981) found that the implicit discount rate over longer time horizons was lower than over shorter horizons, implying time inconsistency without quantifying the magnitude of β (nor specifically supporting a quasi-hyperbolic model over other alternatives). Similar evidence is also well documented in early papers in psychology (e.g., Ainslie and Haendel, 1983).

Despite decades of subsequent work and many dozens of studies, there is little consensus on how to best measure discounting (Cohen et al., 2020). Recently, several notable studies that control carefully for confounding factors in the elicitation procedure (such as transaction costs and trust in the researcher) find no evidence of present bias for monetary rewards (Andersen et al., 2014; Andreoni and Sprenger, 2012; Augenblick et al., 2015). As a result, it is becoming a new stylised fact that present bias either does not exist for money, or that it is an artefact of the study design and procedures.

Given the popularity of quasi-hyperbolic discounting in both applied and theoretical work across multiple disciplines, it is important to consolidate and evaluate the strength of evidence for it in the literature, as well as to examine the sources of heterogeneity in both present bias and long-run impatience. In this paper, we do this by means of a systematic meta-analysis, drawing on a comprehensive search from

which we constructed a database of 108 estimates of the present bias parameter β (88 for money and 20 for other rewards) from 86 studies, and 94 matched estimates of the (annualised) discount factor δ .

Our meta-analytic estimate of β for money is 0.919 with 95% confidence interval [0.891, 0.948], indicating a modest but significant present bias. However, we detect selective reporting in favour of smaller estimates of β that imply a stronger present bias. After correcting for selective reporting, our most conservative estimate of β for money (in the sense of implying the least present bias) is 0.938, with 95% confidence interval [0.905, 0.972]. For non-monetary rewards, our meta-analytic estimate of β is 0.750, with 95% confidence interval [0.643, 0.857], and there is no evidence of selective reporting. Disaggregating to specific non-monetary rewards, we find that present bias for real effort (implied β of 0.815 and 95% confidence interval [0.725, 0.904]) is intermediate between that for money and for other rewards such as real consumption.

After controlling for β , our meta-analytic estimate of δ for money is 0.745 with 95% confidence interval [0.680, 0.809], corresponding to an annual discount rate of 34.27% [22.65%, 45.90%]. For non-monetary rewards it is 0.786 with 95% confidence interval [0.587, 0.985], corresponding to an annual discount rate of 27.24% [-5.03%, 59.52%].

We find there is considerable heterogeneity in estimates of β and δ , but we have only modest success in explaining it. Of particular note, we find that studies that correct for potential non-linearity of the utility function are associated with larger estimates of β (weaker present bias), but not of delta; the same is true specifically for the Convex Time Budget protocol (CTB, [Andreoni and Sprenger, 2012](#)). Studies that have a larger median interest rate find more present bias but not more impatience. Intriguingly, we find no effect of real versus hypothetical incentives on either β or δ . Finally, we find that Asian samples tend to be more present biased, while European ones are more patient.

Our paper generates novel insights beyond what can be learned from the two existing related meta-analyses. [Imai et al. \(2021\)](#) report a meta-analysis of β focusing *exclusively* on studies that use the CTB protocol. Because our sample is not restricted in this way, we include estimates from more than three times as many studies. Moreover, we are able to compare the results of different elicitation procedures, and confirm that the CTB method is associated with findings of weaker present bias. We also provide an analysis of the matched estimates of δ , after the effects of β are accounted for. [Matousek et al. \(2022\)](#) report a meta-analysis of *discount rates*, pooling estimates of three different discount functions: exponential, hyperbolic, and quasi-hyperbolic. To compare estimates of these distinct functions, [Matousek et al. \(2022\)](#) convert all estimates into an equivalent exponential discount rate, treating estimates of non-exponential models as if they were generated by a true exponential process. In this way, they quantify the *magnitude* of discounting while abstracting from its *source*. By focusing specifically on estimates of the quasi-hyperbolic model, our meta-analysis can distinguish between the two sources of discounting recognised by that model, namely present bias and long-run impatience.

Nonetheless, our meta-analysis is constrained by the limitations of working with the available secondary data. Not only do we lack access to most primary datasets, we also cannot recover the correlations between the estimates of β and δ for the vast majority of studies. As a result, we are unable to conduct a fully joint meta-analysis of β and δ that would account for their structural interdependence in the quasi-

hyperbolic model. Instead, we report the results of independent univariate meta-analyses of β and δ , separately also for monetary and non-monetary rewards.

The remainder of the paper proceeds as follows: Section 2 sets out our methodology, Section 3 presents our results, and Section 4 provides a discussion.

2 Methodology

2.1 Conceptual framework

The classical model of exponentially discounted utility (Samuelson, 1937) assumes that a decision-maker’s intertemporal preferences are governed by a parameter δ , called the discount factor, and that when making a plan today she attaches a weight δ^t to the subjective value of payoffs t periods in the future. A key implication of exponential discounting is time consistency (Strotz, 1955): optimal plans do not change as a result of the mere passing of time. The quasi-hyperbolic β - δ model (Laibson, 1997) applies an extra discount β to all future rewards ($t > 0$) to capture time inconsistency in the form of present bias.¹ In the β - δ model, a decision-maker (at time 0) values a stream of payoffs $(x_0, \dots, x_T)$ as:

$$U(x_0, \dots, x_T) = u(x_0) + \beta \sum_{t=1}^T \delta^t u(x_t)$$

where $0 < \delta < 1$ captures impatience (while $\delta > 1$ would indicate negative discounting), $0 < \beta < 1$ captures present bias (while $\beta > 1$ would indicate future bias), and $u(x_t)$ is instantaneous utility at time t . When $\beta = 1$, the β - δ model reduces to the standard exponential model.

2.2 Dataset construction

Appendix A and Appendix Figure A.1 document our strategy to identify relevant studies for the meta-analysis. Because we focus on the β - δ model, we restrict attention to studies that report estimates of β . Where available, we match each estimate of β with its corresponding estimate of δ , converted (where necessary) into annualised terms. The full list of 86 included studies is reported in Appendix B.

In addition to point estimates of β and δ , we require information on their precision in the form of standard errors, for the purpose of weighting the estimates in a meta-analysis. Studies differ in how they report their results. For example, some provide aggregate estimates while others provide summary statistics of individual estimates. Where individual-level data are reported, we transform the standard deviation of the individual estimates into the standard error of the mean estimate. Where standard errors are not reported directly, we reconstruct them from other available information such as t -ratios or p -values.

¹Strictly speaking, a finding that $\beta < 1$ can be interpreted as evidence of dynamic inconsistency only if it is inferred from a *longitudinal design* in which participants face identical temporal trade-offs at two distinct points in time: one in which the sooner outcome is in the future, and another in which it is available immediately. In practice, only a minority of the estimates in our dataset (24 out of a total of 108 estimates of β) are derived from such designs. The remainder identify β by varying the front-end delay across choices made at a single point in time. In the latter design, a finding that $\beta < 1$ provides direct evidence only that preferences are non-stationary and does not automatically imply time inconsistency except under the additional assumption of temporal stability (Halevy, 2015).

Studies also vary in how they report information on δ : while the majority explicitly report discount factors, others instead report a discount rate ρ , where $\delta = 1/(1+\rho)$. They also differ in the time interval used to compute δ or ρ (e.g., daily, weekly, monthly, or annual). We convert discount rates into discount factors, and express all discount factors in annualised terms, applying the delta method to compute standard errors where discount factors had to be recalculated.²

Some studies report multiple estimates of β and δ . Where these represent a full sample and its subsamples (e.g., males and females), we record one estimate for the full sample and do not include estimates for the subsamples. Where a study reports multiple estimates derived from a single dataset, we record the estimate reported as the main result (omitting sensitivity analyses and robustness checks). These procedures minimise interdependence resulting from the inclusion of multiple estimates from the same dataset. However, where a study reports multiple estimates as a result of collecting multiple datasets from a single sample (for example, when comparing different elicitation methods or reward types within subjects),³ we include all of these estimates. This allows us to investigate how the choice of elicitation procedure or reward type affects the resulting estimates of β and δ .

Through these procedures, the 86 studies yield 108 estimates of β (88 for money and 20 for other rewards). From the same studies, we obtain 94 matched estimates of δ (79 for money and 15 for other rewards). We have fewer estimates of δ than β because some studies either do not estimate δ or assume it to be equal to one. We have six estimates of β , and ten of δ , for which standard errors could not be recovered. We adopt the procedure in [Brown et al. \(2024\)](#) to impute missing standard errors, and in Appendix Table D.3 we verify that we obtain very similar results when these observations are instead excluded.

To investigate sources of heterogeneity in estimates of β and δ , our dataset captures a wide range of methodological features of the studies and estimates, including subject pool, incentives, elicitation method, whether and how a study measures utility curvature, interest rates and stake sizes, estimation method, immediacy of sooner rewards, location, and discipline. See Appendix C for details of these variables.

2.3 Meta-analytic methods

In this section we document our statistical methods. The reader can skip ahead to the results in Section 3 with minimal loss of continuity.

To synthesise the body of evidence on quasi-hyperbolic discounting, we report the results of a random-effects (RE) meta-analysis model, as well as a multilevel random-intercepts model that allows for potential correlation of estimates within author-group clusters. We estimate these models separately for β and δ , and for monetary and non-monetary rewards. We report these models using our full dataset, including estimates for which we had to impute standard errors; in the Appendix we report very similar results using only estimates for which standard errors were recoverable. We then test for selective reporting,

²If a study reports an n -period discount rate (e.g., a daily discount rate ρ_n , where n is the number of compounding periods per year), the annualised discount factor is $\delta = 1/(1+\rho_n)^n$. If a study reports an n -period discount factor δ_n (e.g., a monthly discount factor), the annualised discount factor is $\delta = (\delta_n)^n$.

³[Andreoni et al. \(2015\)](#) is an example of a study that compares elicitation methods within subjects, while [Augenblick et al. \(2015\)](#) is an example of a study that compares reward types.

and report results of a broad range of strategies to correct for possible reporting biases. Finally, we report meta-regression analyses to explore sources of heterogeneity in estimates of β and δ .

We acknowledge at the outset that our approach of estimating separate univariate meta-analyses for β and δ reflects the constraints of working with the available secondary data, and cannot account for the structural interdependence of the parameters in the quasi-hyperbolic model. Given complete individual data from each primary study, one could in principle pool the data from all studies to jointly estimate β and δ in the grand dataset of individual primary data. Where only secondary data is available, a multivariate meta-analysis (MVMA) model (Jackson et al., 2020) can be used to jointly synthesise multiple outcomes of interest while accounting for their correlation. Unfortunately, MVMA requires that each study report the within-study covariances or correlations of the outcome variables, which is seldom reported in the primary literature. Using the 21 estimates (out of a total of 108) for which we were able to recover the necessary correlations between β and δ , we found the results of univariate and multivariate analyses to be almost identical; we believe this reflects the fact that the variance in our data is dominated by between-study variance (see Section 3.6), such that within-study correlations have little impact. In light of this evidence, and given that the MVMA model would necessitate discarding the bulk of our data, we report the results of separate univariate models for β and δ . We also concede that we cannot assess the extent to which these models are biased relative to the hypothetical benchmark of a structural model in which β and δ are jointly estimated from the pooled individual primary data.

2.3.1 Baseline models

Letting $\theta \in \{\beta, \delta\}$ denote the model parameter under consideration, the RE model (Hedges, 1983; DerSimonian and Laird, 1986) treats each observation $\hat{\theta}_j$ as an estimate of a different unknown true θ_j , where the $\{\theta_j\}$ are assumed to be a random sample from a larger population of interest:

$$\hat{\theta}_j = \theta_j + \varepsilon_j = \theta + u_j + \varepsilon_j \quad (1)$$

where $\varepsilon_j \sim N(0, \sigma_j^2)$ represents sampling error, $u_j \sim N(0, \tau^2)$ captures between-observation heterogeneity beyond sampling variance, and ε_j and u_j are assumed to be independent.

The objective is to estimate $\theta = E(\theta_j)$, the mean of the underlying distribution of θ_j . The random-effects estimate $\hat{\theta}_{RE}$ is a weighted average of the individual $\hat{\theta}_j$:

$$\hat{\theta}_{RE} = \frac{\sum_j w_j \hat{\theta}_j}{\sum_j w_j}$$

where the weights are given by $w_j = 1 / (\hat{\sigma}_j^2 + \hat{\tau}^2)$, and are thus inversely related to the total (within-plus between-observation) variance, such that estimates that have greater precision (smaller standard errors $\hat{\sigma}_j$) are given greater weight. The between-observation variance τ^2 is estimated using the method of restricted maximum likelihood (Raudenbush, 2009).

Our dataset sometimes includes multiple studies by the same authors: from the 86 studies we identify 56 distinct author clusters with no overlapping co-authors. To account for the possibility that estimates

from the same cluster may be correlated, we estimate a multilevel extension of the RE model (Goldstein et al., 2000; Thompson et al., 2001). The three-level (3L) random-intercepts model is given by:

$$\hat{\theta}_{ij} = \theta + v_i + u_{ij} + \varepsilon_{ij} \quad (2)$$

where $\hat{\theta}_{ij}$ is the j -th estimate within the i -th cluster, and $v_i \sim N(0, \tau_v^2)$, $u_{ij} \sim N(0, \tau_u^2)$ and $\varepsilon_{ij} \sim N(0, \sigma_{ij}^2)$ are assumed independent. In the paper we report a 3L model using author groups as the clustering variable. We also estimate a version using study clusters, and a four-level (4L) model with estimates nested within studies nested within author groups. Since the results are similar to those using author group clusters, with substantially overlapping confidence intervals, and since the Bayesian Information Criterion (BIC) favours the version with author-group clusters, we do not report those models.

2.3.2 Selective reporting

Selective reporting refers to systematic biases in the reporting of findings based on their magnitude, direction, or statistical significance, which can result in a literature overstating the weight of evidence supporting the predictions of a dominant model (for example, that $\beta < 1$ in the quasi-hyperbolic model). The starting point in the analysis of selective reporting is the classic Egger equation (Egger et al., 1997), which is used to test for the presence of unexpected correlation between the reported estimates $\hat{\theta}_j$ and their standard errors $\hat{\sigma}_j$ in the following (weighted) linear regression:

$$\hat{\theta}_j = \alpha_0 + \alpha_1 \cdot \hat{\sigma}_j + \varepsilon_j \quad (3)$$

In the absence of selective reporting, there should not be any systematic relation between the reported estimates and their precision. The test of the null hypothesis that $\alpha_1 = 0$ in equation (3) is known as the Egger test, or funnel asymmetry test. Since the Egger test is seen as having low power, it is considered best practice to report estimates that correct for possible selective reporting even where none is detected according to the test (Stanley and Doucouliagos, 2014; Irsova et al., 2024). We report results of several such corrections, and we describe the logic of these procedures next.

The PEESE correction (Stanley and Doucouliagos, 2014) builds on the intuition that estimates with very small standard errors are unlikely to be selectively reported, as they are almost sure to detect a significant effect: any systematic relation with the standard error would only be evident in less precise estimates. To capture this non-linear effect of precision, the PEESE model replaces the standard error on the right-hand side of equation (3) with the variance, to estimate the following regression:

$$\hat{\theta}_j = \alpha_0 + \alpha_1 \cdot \hat{\sigma}_j^2 + \varepsilon_j \quad (4)$$

and uses the estimate of the intercept in equation (4) as the corrected measure of the true effect beyond bias as standard errors approach zero.

The endogenous kink (EK) model (Bom and Rachinger, 2019) similarly builds on the insight that sufficiently precise estimates are unlikely to be selectively reported, but replaces the quadratic term in (4) with a piecewise-linear function that takes the value zero up to a cut-off level of $\hat{\sigma}_j$ below which selective reporting is unlikely. To locate this cut-off, EK compares a first-stage estimate to an assumed 5% threshold of statistical significance. The intercept of the (weighted) piecewise-linear regression is the corrected estimate of θ . Should the cut-off value of $\hat{\sigma}_j$ fall at zero (which occurs when the aggregate effect is small or imprecisely estimated), EK reduces to the estimate of the intercept in equation (3).

The WAAP estimate (“weighted average of the adequately powered”, Ioannidis et al., 2017) is a weighted average computed using only those estimates that have adequate power to discriminate the initial random-effects estimate from one. Assuming a 5% significance level and 80% threshold for power, a study has adequate power to discriminate $\hat{\theta}_{RE}$ from one if its standard error is less than $(1 - \hat{\theta}_{RE}) / 2.8$.⁴ The WAAP is obtained simply by re-estimating the model using only those estimates that pass this test.

Finally, the MAIVE estimator (Irsova et al., 2023) corrects for spurious precision (which operates through selection on smaller standard errors rather than larger effect sizes) by using the inverse sample size as an instrument for the variance in the PEESE model of equation 4.

2.3.3 Heterogeneity analysis

To explore the sources of heterogeneity in estimates of θ , we pursue two approaches. First, we run a series of meta-regressions, each of which examines the effect of including moderators for a different dimension of heterogeneity. This analysis faces the challenge that not all of the variables are equally important: some are redundant, and including them may reduce the precision of the point estimates for more important variables. To address this model uncertainty, our second approach is to employ Bayesian Model Averaging (BMA, Raftery et al., 1997; Steel, 2020) to identify the variables that are most likely to influence θ . BMA estimates many regression models spanning the entire space of possible combinations of the explanatory variables. From this, it computes a weighted average of the estimated coefficients, leveraging posterior model probabilities (PMPs) obtained using Bayes’ theorem.⁵ BMA also computes posterior inclusion probabilities (PIPs) for each variable, reflecting the cumulative posterior model probability of all models in which that variable is included.

3 Results

3.1 Characteristics of studies and estimates

Appendix Table D.1 reports characteristics of the 86 studies in our dataset. 85% of the studies originate from the fields of economics and business, reflecting the fact that the β - δ model is the dominant behavioural theory of discounting in these disciplines. 74% of studies report estimates from developed countries, and 81% were incentivised. As of February 2025, 27% of the studies were unpublished.

⁴In the funnel plots reported in Figure 2, these are estimates whose standard errors are smaller than the threshold depicted by the horizontal dashed line.

⁵Our BMA analysis uses a dilution model prior (George, 2010) to account for collinearity across variables, and the unit information g-prior recommended by Eicher et al. (2011).

Appendix Table D.2 reports characteristics of the 108 estimates of β (recall that our 94 matched estimates of δ correspond to a subset of these). 41% were derived from students, and a further 44% from general adult samples; the remainder study specialised populations such as adolescents and clinical samples. The choice list (Coller and Williams, 1999) is the most common elicitation method, accounting for 44% of estimates, including ones obtained by joint elicitation (Andersen et al., 2008). 43% of estimates are obtained using the Convex Time Budget (CTB) method of Andreoni and Sprenger (2012). Appendix Figure D.1 illustrates how the prevalence of these methods has evolved over time. 63% of estimates control for utility curvature; this includes both joint elicitation and CTB methods. Appendix Figure D.2 reports a plot of each estimate of β against its year of publication in panel A, and year of data collection in panel B. In this figure, different markers denote different elicitation methods.

3.2 Present bias for money

The first row in Table 1 reports summary statistics of the 88 estimates of β for money, while panel A in Figure 1 reports a histogram. The mean estimate is 0.910, and its 95% confidence interval is [0.876, 0.943]. The distribution of estimates is left-skewed, such that the median of 0.964 implies less present bias than the mean. 42 estimates (48%) are significantly smaller than one indicating present bias, eight (9%) are significantly greater than one, 32 (36%) are not significantly different from one, and the remaining six (7%) have missing standard errors.

The mean estimate reported above does not consider the precision of the estimates. To account for this, we use the random-effects model in equation (1) to establish a meta-analytic weighted average estimate of β , after first imputing missing standard errors for six observations.⁶ The resulting estimate is reported in the first cell in the first row of Table 2: the random-effects estimate of β for monetary rewards is 0.919 with a 95% confidence interval of [0.891, 0.948]. This is significantly less than one, indicating present bias for money in the β - δ model.

Appendix Figure E.1 reports the forest plot (Hedges and Olkin, 1985) of the individual estimates of β for monetary rewards. Each row corresponds to a different estimate of β . The size of each box represents the weight of that observation in calculating the random-effects estimate, and the horizontal line around the box represents the 95% confidence interval of that estimate. The overall random-effects estimate is depicted by the diamond at the bottom of the figure.

To account for the fact that our dataset includes multiple estimates of β from overlapping author groups, the second cell in the first row of Table 2 reports a three-level random intercepts model (equation 2) which generalises the random-effects model to allow for estimates from a given author cluster to be correlated. The point estimate of β decreases slightly to 0.915, but its standard error also increases such that the 95% confidence interval of [0.877, 0.953] fully contains that of the random-effects estimate.

Selective reporting of β may occur if certain findings are more likely to be reported (and published) depending on their sign, magnitude, or statistical significance. Thus if estimates of $\beta < 1$ are favoured by selective reporting, this would bias the body of evidence toward the prediction of the dominant theory, and exaggerate the strength of present bias. Alternatively, contrarian findings of β close to one may be

⁶Appendix Table D.3 reports estimates that exclude observations with missing standard errors, which are very similar to those described in the text.

Table 1: Summary statistics of β and δ estimates

	N	Mean	SD	P10	P25	P50	P75	P90
β , Money	88	0.910	0.159	0.674	0.882	0.964	0.995	1.020
β , Non-money	20	0.754	0.241	0.347	0.610	0.832	0.929	0.987
δ , Money	79	1.583	7.490	0.173	0.596	0.832	0.996	1.001
δ , Non-money	15	78.156	165.758	0.740	0.868	0.949	6.175	381.510

Figure 1: Distribution of β and δ estimates




Table 2: Meta-analytic estimates of β and δ

	RE	3L-authors	PEESE	EK	WAAP	MAIVE
β , Money	0.919 (0.014) [0.891, 0.948]	0.915 (0.019) [0.877, 0.953]	0.924 (0.015) [0.895, 0.953]	0.938 (0.017) [0.905, 0.972]	0.936 (0.013) [0.910, 0.962]	0.809 (0.174) [0.468, 1.150]
N	88	88	88	88	64	88
β , Non-money	0.750 (0.055) [0.643, 0.857]	0.735 (0.070) [0.597, 0.873]	0.745 (0.058) [0.631, 0.859]	0.750 (0.082) [0.589, 0.912]	0.747 (0.058) [0.634, 0.860]	0.755 (0.099) [0.560, 0.950]
N	20	20	20	20	16	20
δ , Money	0.745 (0.033) [0.680, 0.809]	0.710 (0.051) [0.610, 0.810]	0.745 (0.033) [0.680, 0.809]	0.763 (0.036) [0.693, 0.833]	0.757 (0.034) [0.690, 0.824]	0.749 (0.053) [0.644, 0.853]
N	79	79	79	79	70	79
δ , Non-money	0.786 (0.102) [0.587, 0.985]	0.771 (0.115) [0.545, 0.998]	0.786 (0.102) [0.586, 0.985]	0.762 (0.102) [0.562, 0.962]	0.786 (0.115) [0.561, 1.012]	0.465 (0.093) [0.283, 0.646]
N	15	15	15	15	8	15

Notes: Standard errors are in parentheses, and 95% confidence intervals are in square brackets. The random effects (RE) model treats each observation as an estimate of a different true effect, and estimates the mean of the distribution of these effects. The 3L-authors model adds random intercepts to allow for the possibility that estimates originating from the same author-group cluster may be correlated. PEESE, EK, WAAP, and MAIVE are alternative procedures to correct meta-analytic estimates for the possibility of reporting biases, as described Section 2.3.2; in particular, WAAP is based on a restricted sample of “adequately powered” estimates.

reported as support for the normative model of exponential discounting. However, findings of $\beta > 1$ may be viewed with suspicion and thus discarded.

To formally test for selective reporting, we estimate the Egger equation (equation 3), a regression of each estimate of β on its standard error (Egger et al., 1997). The Egger test is based on the idea that, in the absence of selective reporting, there should be no systematic relation between the estimates and their precision. We instead find that estimates with larger standard errors tend to be associated with smaller values of β ($z = -2.053$, $p = 0.040$), indicating selective reporting in favour of smaller estimates of β .

Figure 2 panel A reports the funnel plot (Light and Pillemer, 1984) of the estimates of β for monetary rewards. This is a scatter plot of the estimates on the horizontal axis against their standard errors on the vertical (with the scale reversed, such that more precise estimates appear toward the top). The light grey vertical line corresponds to the benchmark value of $\beta = 1$. The dark vertical line depicts the random-effects estimate, while the diagonal arms represent its 95% pseudo confidence interval: estimates falling within these arms do not differ significantly from the random-effects estimate at the 5% level. The dashed horizontal line is the threshold value of the standard error at which an estimate has 80% power to discriminate the random-effects estimate from one: 64 (73%) of the 88 estimates pass this test.

The asymmetry of the funnel plot suggests there may be “missing studies”: there are more observations to the bottom left compared to the bottom right, indicating that estimates of β that are greater than the mean and have large standard errors are less likely to be reported. The diagonal dashed red line is the fitted Egger regression: in the absence of selective reporting, this would be vertical. Consistent with the finding of the Egger test, we instead see that the fitted line is upward sloping: less precise estimates toward the bottom of the funnel tend to be associated with smaller estimates of β .

The remaining cells in the first row of Table 2 report a diverse range of strategies to adjust the meta-analytic estimate for the possibility of reporting biases. The logic behind each of these procedures is described in Section 2.3.2. Consistent with the finding of selective reporting of β for money, the results generally indicate a moderate attenuation of β toward one.⁷ The WAAP procedure uses the restricted sample of 64 adequately-powered estimates to obtain an estimate of 0.936, with 95% confidence interval [0.910, 0.962]. The most conservative estimate, in the sense of indicating the least present bias, is given by the endogenous kink procedure, which gives an estimate of β for monetary rewards of 0.938 with 95% confidence interval [0.905, 0.972], which remains significantly smaller than one.⁸


3.3 Present bias for non-monetary rewards

The second row in Table 1 reports summary statistics of the 20 estimates of β for non-monetary rewards, while panel B in Figure 1 reports the corresponding histogram. Of these estimates, 12 are based on real effort as the (aversive) reward (following Augenblick et al., 2015), 4 use food as a real consumption reward, and the remaining 4 involve other rewards such as health or environmental outcomes. The mean estimate is 0.754, and its 95% confidence interval is [0.642, 0.867]. The distribution is again left-skewed, with the median of 0.832 indicating less present bias than the mean. 15 of the estimates (75%) are significantly smaller than one, while the remaining 5 (25%) do not differ significantly from one.

⁷The exception is MAIVE, which has a poor precision reflecting the fact that the inverse sample size turns out to be a weak instrument for the variance.

⁸Since the kink is found to be located at zero, this is equivalent to the estimate of the intercept of the Egger equation (3).

Figure 2: Funnel plots of β and δ estimates



Notes: In each panel, the dark vertical line depicts the random-effects meta-analytic estimate, while the grey vertical line corresponds to the benchmark value of one. Estimates falling within the diagonal arms are consistent with the meta-analytic estimate at the 5% significance level. The dashed horizontal line represents the threshold value of the standard error at which an estimate is adequately powered to discriminate an effect of the size implied by the meta-analytic estimate from one. The red dashed line is the fitted Egger regression. Panel C omits one estimate with a weight of 0.00%; Panel D omits six estimates with an aggregate weight of 0.76%.

The random-effects estimate of β for non-monetary rewards in Table 2 is 0.750 with 95% confidence interval [0.643, 0.857]. This is significantly smaller than one, indicating that there is present bias for non-monetary rewards. The confidence interval is wider than for money since there are fewer estimates for non-monetary rewards; nonetheless, the two confidence intervals do not overlap. The forest plot of the estimates is reported in Appendix Figure E.2.

The Egger test does not detect selective reporting of β for non-monetary rewards ($z = -0.009$, $p = 0.993$). In line with this, the funnel plot in Figure 2 panel B is symmetric, and the fitted Egger regression is vertical and coincides with the random-effects estimate. Also consistent with the finding of no selective reporting, the results of all four correction procedures reported in Table 2 yield results very close to the random-effects estimate, with point estimates varying between 0.745 and 0.755.

Thus, consistent with a widely held view in the literature, we find stronger present bias for non-monetary compared to monetary rewards. To examine the effect of specific reward types, we run a random effects meta-regression pooling estimates for both money and non-monetary rewards, with dummy variables for each reward type as the regressors (and money as the omitted category). We find that the coefficient for every non-monetary reward is significantly negative (for real effort $p = 0.031$, for food $p = 0.003$, and for other rewards $p < 0.001$), indicating stronger present bias than for money. The coefficients for food and for other rewards do not differ significantly from one another, with $p = 0.856$ in a Wald test. For real effort the implied estimate of β is 0.815 with 95% confidence interval [0.725, 0.904], for food it is 0.656 [0.487, 0.825], and for other rewards 0.634 [0.478, 0.791]. Thus, present bias for real effort is more pronounced than for money, but weaker than for other non-monetary rewards.

3.4 Impatience for monetary rewards

The third row in Table 1 reports summary statistics for 79 estimates of δ for money, while panel C of Figure 1 reports a histogram. The unexpectedly large mean and standard deviation reflect the presence of a single outlier, which is the result of annualising a weekly estimate of δ larger than one. Because annualising also magnifies the standard error, this estimate has negligible weight (less than 0.001%) in the random effects model. Omitting this observation, the unweighted mean of δ is 0.741 with 95% confidence interval [0.673, 0.808], corresponding to an annual discount rate of 35.03%. The median estimate is 0.832, which represents an annual discount rate of 20.12%. 53 estimates (67%) are significantly smaller than one, 3 (4%) are significantly greater than one (indicating negative discounting), 14 (18%) are not significantly different from one, and the remaining 9 (11%) have missing standard errors.

Accounting for the precision of the estimates, the random-effects estimate of δ for money in Table 2 is 0.745 with 95% confidence interval [0.680, 0.809], corresponding to an annual discount rate of 34.27% [22.65%, 45.90%]. The forest plot is in Appendix Figure E.3. The Egger test does not detect selective reporting of δ for money ($z = -1.186$, $p = 0.236$), although the fitted Egger equation in the funnel plot in Figure 2 panel C does have a discernible upward slope.⁹ Accordingly, the effect of applying corrections for reporting biases is to slightly attenuate δ toward one. 70 (89%) of the 79 estimates have adequate

⁹We omit the single outlier from both the forest plot in Appendix Figure E.3 and the funnel plot in Figure 2 panel C to improve the legibility of these figures. However, we include this observation in both the random-effects model (which also determines the threshold for “adequate power”) and in the estimated Egger regression.

power to discriminate the random-effects estimate from one. Using this sample, the WAAP estimate is 0.757 [0.690, 0.824], and corresponds to an annual discount rate of 32.16% [20.46%, 43.86%].

3.5 Impatience for non-monetary rewards

Our results for δ for non-monetary rewards should be interpreted with caution for several reasons. First, out of 20 estimates of β for non-monetary rewards, we have matched estimates of δ for only 15, indicating that the proportion of studies that do not estimate δ and instead assume it to be one is greater than for monetary rewards. Second, studies of non-monetary rewards may differ also in an important feature of their design. In particular, whereas it is generally not common for studies of discounting to allow for the possibility of $\delta > 1$ or negative discounting (Lipman and Attema, 2024), it is more common to do so in studies of non-monetary rewards, especially in studies of real effort following Augenblick et al. (2015), but also in our own study of dietary rewards (Cheung et al., 2022). Where $\delta > 1$ is found, annualisation can result in both the estimate and its standard error being greatly magnified.

As can be seen in both the summary statistics in Table 1 and histogram in panel D of Figure 1, our limited sample of only 15 estimates also includes several very large estimates of $\delta > 1$. In the funnel plot in Figure 2 panel D we omit six observations (five with large point estimates of δ , and another with a large standard error) for ease of visibility of the remaining data, yet these six observations have a *combined* weight of only 0.76% in the random-effects estimate. This means that our results for δ for non-monetary rewards are driven almost entirely by the remaining data points displayed in the figure.

With these caveats in hand, the 15 estimates comprise 9 for real effort, 4 for food, and 2 for other rewards. 8 are significantly less than one, 6 do not differ significantly from one, and the remaining estimate has a missing standard error. The random effects estimate of δ for non-monetary rewards is 0.786 with 95% confidence interval [0.587, 0.985], corresponding to an annual discount rate of 27.24% [−5.03%, 59.52%]. The forest plot of the estimates is reported in Appendix Figure E.4. The Egger test detects marginal evidence of selective reporting ($z = 1.770$, $p = 0.077$) in favour of larger estimates of δ , although this is not borne out by the results of the correction procedures.

3.6 Heterogeneity analysis

In this section, we summarise the results of our analysis of the sources of heterogeneity in estimates of β and δ . We focus on monetary rewards because we have insufficient observations to perform this analysis for non-monetary rewards. The I^2 statistic (Higgins and Thompson, 2002) quantifies the proportion of the total variation in the estimates that is due to between-study heterogeneity, as opposed to sampling error. We find that essentially all variance is due to true differences in the underlying effects, with $I^2 = 99.95\%$ for β , and $I^2 = 100\%$ for δ . To explain this heterogeneity, we coded information on a wide range of characteristics of the studies and estimates, as detailed in Appendix C. We describe two approaches to examining the effects of these characteristics, and we report the full results in Appendix F and G.

First, in Appendix Table F.1 for β , and Appendix Table F.2 for δ , we report a series of meta-regression models, where each model focuses on the effects of a single dimension of heterogeneity (for example,

subject pools in model 2).¹⁰ Second, to address the issue of model uncertainty we use Bayesian Model Averaging (BMA), estimating many models spanning the space of possible combinations of regressors to identify the variables that are most likely to influence β and δ .

Results of the BMA analysis are visualised in Appendix Figure G.1 for β , and Appendix Figure G.2 for δ . In these figures, the explanatory variables are displayed on the vertical axis and sorted in descending order by posterior inclusion probability (PIP, the sum of the model probabilities of those models that include the variable), such that more influential variables appear toward the top. On the horizontal axis, the width of each column represents the posterior model probability (PMP) of a model, and the models are sorted in descending order by PMP. The shading of each cell indicates the sign of the coefficient of the corresponding variable in the respective regression, with darker grey indicating a positive sign and lighter grey a negative sign. Blank cells indicate that a variable is not included in a given model. Numerical results of the BMA analysis are provided in Appendix Table G.1 for β , and Appendix Table G.2 for δ , reporting the posterior mean and standard deviation for each variable, together with its PIP.

In the meta-regression analyses for β in Appendix Table F.1, model 5 which focuses on methodological variables finds that estimates that control for non-linear utility tend to be associated with larger values of β (less present bias). This variable captures both 21 estimates from 20 studies that adopt the joint elicitation approach following [Andersen et al. \(2008\)](#), as well as 34 estimates from 24 studies that adopt the CTB design of [Andreoni and Sprenger \(2012\)](#). The same model finds that estimates obtained using hypothetical rewards (20 estimates from 16 studies) do not differ significantly from ones that use real incentives: the point estimate of this effect is 0.016, and its t -statistic is 0.33. In model 6 which focuses on elicitation procedures (with choice list as the omitted category), estimates derived from the CTB procedure are associated with less present bias, while those in the residual category (e.g., BDM auction) are associated with stronger present bias. In model 7 which focuses on design parameters, estimates obtained from designs that have a larger median interest rate are associated with a stronger present bias. In the grand regression model that includes all of the covariates, the effects of CTB elicitation (weaker present bias) and median interest rate (stronger present bias) remain significant at the 5% level.

In the BMA analysis for β in Appendix Table G.1, we find that estimates from Asian samples (18 estimates from 16 studies) exhibit stronger present bias (with a PIP of 55%), while those that use the CTB design detect less present bias (with a PIP of 40%). In Appendix Figure G.1, the four most probable models, with a combined PMP of 44%, are the four permutations obtained by including or excluding these two variables alone. The next most influential variable, with a PIP of 19%, identifies studies conducted in schools or workplaces (6 estimates from 5 studies), which tend to find a stronger present bias. All remaining variables have PIPs of 10% or less.

Our heterogeneity analyses for δ omit the single outlier noted in Section 3.4. The meta-regressions indicate that studies from Europe (22 estimates from 19 studies) find greater patience compared to the omitted category of North America (model 4), and that studies that use “other” elicitation procedures are associated with greater patience (model 6), as are studies that have a larger median interest rate (model

¹⁰We report these meta-regressions with standard errors clustered on author groups. To account for the limited number of clusters, in Appendix Table F.3 for β , and Appendix Table F.4 for δ , we additionally report confidence intervals obtained by wild bootstrapping.

7). The effects of the latter two variables go in the opposite direction to their effects upon β . Of these variables, only Europe remains significant at the 5% level in the grand regression. In contrast to the corresponding analysis for β , model 5 indicates that controlling for non-linear utility has no effect upon estimates of δ ; in the same model, it is again the case that there is no effect of hypothetical rewards. In addition, the grand regression indicates that, after controlling for all other variables, studies from outside the disciplines of economics and business tend to find less patience, at the 1% level of significance.

In the BMA analysis for δ , Appendix Figure G.2 indicates that the most probable model, with a PMP of 52%, is the null model that includes only an intercept. Appendix Table G.2 indicates that the two most influential variables are Europe (with a PIP of 21%), and OLS estimation (with a PIP of 17%). All remaining variables have PIPs of 10% or less.

4 Discussion

Both in business and in private life, managers and decision-makers consistently fail to follow through on their plans, especially where those plans entail immediate costs and delayed benefits. They pledge to be more financially responsible and to prioritise long-term goals such as research and development, but also to exercise more, eat more healthily, or to quit smoking – only to fail to execute these plans, often to their own frustration and disappointment. In behavioural economics, the dominant model of such self-control problems is quasi-hyperbolic discounting. The central tenet of this model is that decision-makers have a “present-bias” toward immediate rewards, in addition to the long-run discounting of delayed outcomes. Despite the popularity of this model across multiple fields of decision sciences, the empirical literature has not reached any consensus over the extent of present bias or impatience, calling for a scientific re-examination of the existing evidence base.

In this paper, we report a meta-analysis of the two parameters in quasi-hyperbolic discounting, β and δ . Our dataset comprises 108 estimates of present bias (β) from 86 studies (both published and unpublished) and 94 matched estimates of the annual discount factor (δ) from the same set of studies. We report random-effects meta-analysis models that account for the precision of these estimates, separately for money and non-monetary rewards, as well as multi-level models that allow estimates to be correlated within author-group clusters. We also report the results of a diverse range of strategies to correct for the possibility of reporting biases. However, we reiterate that owing to the limitations of the available secondary data we were unable to conduct a multivariate meta-analysis of β and δ , and thus report the results of independent univariate meta-analyses for each of these parameters.

For money, our random-effects meta-analytic estimate of β is 0.919, with 95% confidence interval [0.891, 0.948]. However, we also detect selective reporting in favour of smaller estimates of β that imply a stronger present bias. Our most conservative estimate of β for money after correcting for selective reporting is 0.938, with 95% confidence interval [0.905, 0.972]. For non-monetary rewards, our random-effects estimate of β is 0.750, with 95% confidence interval [0.643, 0.857], and this remains substantively unchanged after correcting for potential reporting biases. When we examine the effects of specific non-monetary reward types, we find that present bias for real effort (with an implied β of 0.815 and 95% confidence interval [0.725, 0.904]) is intermediate between that for money and for other rewards such as real consumption. We thus find that there is significant present bias for all reward types.

From a managerial perspective, there is compelling evidence that people struggle to follow through on work-related goals, especially for repetitive and tedious tasks similar to the ones used in research paradigms. Our estimate of present bias for real effort (0.815) implies that the disutility of effort is perceived as 22.7% greater at the moment of execution than in the planning phase. This implies that even when the long-term goals of managers and employees are aligned, sustaining consistent effort remains a persistent organisational challenge. From a health perspective, our findings underscore how many people fail to adhere to long-term dietary intentions, particularly when it comes to highly palatable but nutritionally poor foods, similar to those commonly used in research studies. Based on our estimate of present bias for food, the immediate utility of consuming such foods jumps by 52.4% at the moment of consumption compared to when people make their plans. This sharp discrepancy aligns with the widespread overconsumption of junk food, and highlights the need for policy interventions that support individuals in bridging the gap between short-term impulses and long-term health goals.

The closest work to our meta-analysis of β is by [Imai et al. \(2021\)](#), who focus on the present-bias parameter estimated *only* using the CTB protocol, resulting in a smaller dataset of 28 studies. They do not find selective reporting of β for money, and find the weighted average of β to be 0.98, differing significantly from one at the 5% level. In our more comprehensive dataset we find that there is selective reporting, but that present bias for money persists even after correcting for it. [Imai et al. \(2021\)](#) similarly find a smaller β (stronger present bias) for real effort than for monetary rewards, however they find selective reporting for real effort. In contrast, our meta-analysis includes a wider range of non-monetary rewards, and we do not find selective reporting of β for these rewards.

After controlling for β , our meta-analytic estimate of δ for money is 0.745 with 95% confidence interval [0.680, 0.809], corresponding to an annual discount rate of 34.27% [22.65%, 45.90%]. For non-monetary rewards it is 0.786 with 95% confidence interval [0.587, 0.985], corresponding to an annual discount rate of 27.24% [-5.03%, 59.52%]. We thus find impatience for both money and non-monetary rewards, although our results for the latter should be interpreted with caution owing to the small sample combined with potential measurement error arising from annualisation of the discount factor.

The closest work to our meta-analysis of δ is the study of discount rates by [Matousek et al. \(2022\)](#). In our paper we focus specifically on estimates of the β - δ model, and thus on estimates of long-run impatience that already account for the effect of present bias. In contrast, [Matousek et al. \(2022\)](#) combine studies of exponential, hyperbolic, and quasi-hyperbolic discounting, converting estimates from each of these functional forms into the equivalent exponential discount rate (see their footnotes 1 and 2). As a result, their measure of the discount rate combines the effects of both present bias and long-run impatience. This is the case both because they include studies that do not account for β , and because for studies that report both β and δ , they collapse the two parameters into a single measure.

While the majority of studies in our dataset consider monetary rewards, it has been argued that money may not appropriately measure time preferences given that the discounted utility model is proposed to describe preferences over consumption as distinct from financial flows (see [Cohen et al. 2020](#) for discussion). Several studies that did not account for present bias have compared discount rates for different reward types, generally finding, first, that there is indeed impatience for money, but also, second, that impatience for non-monetary rewards is stronger ([Estle et al., 2007](#); [Reuben et al., 2010](#); [Ubfal, 2016](#)).

The first study to extend this comparison to the β - δ model was an influential paper by [Augenblick et al. \(2015\)](#), who introduced the use of real effort as an aversive reinforcer in the CTB paradigm, and found stronger present bias for real effort than for money. As noted above, this was confirmed in a meta-analysis of the real effort literature by [Imai et al. \(2021\)](#). We go beyond this by considering present bias for other non-monetary reward types. As a result, we contribute the novel insight that while present bias for real effort is indeed stronger than for money, it is in fact weaker than for other rewards such as real consumption, or health and environmental outcomes.

Our analysis of δ for non-monetary rewards reveals another reason why estimates may differ between money and non-monetary rewards. Whereas only 11 out of 79 (14%) of the estimates of δ for money are derived from designs that allow for the possibility of negative discounting ($\delta > 1$), this is the case for 7 out of 15 (47%) of the estimates of δ for non-monetary rewards. This stark contrast reflects two distinct tendencies within the literature. On one hand, it is generally the case that most designs do not allow negative discount rates to be expressed: in a recent review, [Lipman and Attema \(2024\)](#) find that only one-quarter of designs allow for negative discounting. On the other hand, the seminal paper on real effort by [Augenblick et al. \(2015\)](#) provides for an equal number of positive and negative interest rates, and this design has been widely emulated in subsequent studies of real effort. Similarly, our own study of dietary rewards ([Cheung et al., 2022](#)) also includes both zero and negative interest rates.

Allowing for negative discounting may be particularly impactful when impatience is estimated in conjunction with present bias, as these provide two substitutable channels through which discounting may be expressed. In our dataset, we observe estimates of negative discounting for both real effort and consumption. Since these are reported in weekly terms, and we annualise δ for comparability across studies, they result in some very large outliers. While we are careful not to draw strong conclusions from the results, we also believe it is important to draw attention to the underlying source of this issue.

Correcting for non-linear utility is perhaps the most important recent methodological advance in the empirical study of discounting ([Andersen et al., 2008](#)). We find that studies that adjust for utility curvature report larger estimates of β (less present bias) than ones that do not. Given that the rationale for adjusting for utility curvature ought to apply equally to the estimation of long-run impatience (and was indeed first articulated in that context), it is puzzling that we do not observe any similar effect upon δ . One conjecture, suggested by the findings of recent studies that focus specifically on the utility function ([Abdellaoui et al., 2013](#); [Cheung, 2020](#); see [Cheung, 2016](#), for a discussion) is that correction for utility would not be expected to have a substantial effect if the utility function was in fact close to linear.

Another finding of our heterogeneity analysis that may be surprising to many behavioural economists is that we find no significant effect of whether choices are incentivised or hypothetical, either for present bias or long-run impatience. Both of these findings are in fact consistent with those of a recent large-scale study by [Brañas-Garza et al. \(2023\)](#) that focuses specifically upon the effect of incentives, finding no differences for either β or δ . The meta-analysis by [Matousek et al. \(2022\)](#) similarly finds no effect of incentives upon discount rates that reflect the combined influence of both β and δ .

Our data can be used to shed light on a reformulation of the quasi-hyperbolic model, recently proposed by [Bleichrodt et al. \(2022\)](#). They transform the traditional β - δ parametrisation of the model into δ - τ ,

where τ is defined by $\tau = \frac{\ln \beta}{\ln \delta}$. Since τ is the value that solves $\beta = \delta^\tau$, it has a simple intuition: having a present bias is equivalent to (virtually) adding τ periods of waiting time to a delayed payoff.¹¹ Since our dataset contains matched estimates of β and δ , as well as their standard errors, we can compute τ for the majority of studies, and we use the delta method to derive its standard error. We focus on monetary rewards for this exercise, as we have more observations and fewer concerns over influential outliers.

In our full sample of 79 matched estimates of β and δ for money from 64 studies, we compute a random-effects estimate for τ of 0.261 with 95% confidence interval of [0.103, 0.418]. Since our measure of δ is annualised, this implies that an immediate monetary reward is equivalent to waiting an extra three months for a delayed alternative. Thus the severity of present bias should not be judged from the estimate of β alone, but in light of the associated estimate of δ . While our weighted average β of 0.919 for money in Section 3.2 might suggest only modest present bias, when translated into τ equivalent to three added months of back-end delay it is arguably more consequential. If we restrict attention to estimates that are adequately powered for both β and δ (54 estimates from 43 studies), our estimate of τ is 0.308 with 95% confidence interval [0.109, 0.507]; this is equivalent to 3.7 months of added delay.

Finally, we stress that while our meta-analysis consolidates evidence of time inconsistency within the β - δ framework, it should not be read as endorsement of that model over other alternatives to standard exponential discounting. We note, firstly, that *any* discount function other than the exponential necessarily generates time inconsistency. Secondly, while theoretical work proposes a rich array of alternative models (for example, [Bleichrodt et al., 2009](#); [Ebert and Prelec, 2007](#); [Loewenstein and Prelec, 1992](#); [Read, 2001](#); [Scholten and Read, 2006, 2010](#)), these models are often motivated by behavioural phenomena that are quite distinct from the pattern of present bias for which the β - δ model is invoked.

Thirdly, and most salient to this meta-analysis, in practice the bulk of research seeking to quantify the parameters of specific discount functions has focused on much narrower classes of models. Specifically, research in behavioural economics focuses largely on the exponential and quasi-hyperbolic models, while research in psychology typically estimates either the simple hyperbolic model ([Mazur, 1987](#)) or one of its hyperboloid generalisations ([Myerson and Green, 1995](#); [Rachlin, 2006](#)). The challenge is that these models are identified using quite distinct designs. The β parameter of quasi-hyperbolic discounting is identified by varying the presence or absence of a front-end delay, often accompanied by minimal variation in back-end delays. By contrast, identification of the k parameter of hyperbolic discounting requires extensive variation in back-end delays, but does not rely on any front-end delay. As a result, most datasets in the literature can be used to estimate only one of these two popular classes of models.

Faced with this challenge, the meta-analysis by [Matousek et al. \(2022\)](#) collapses estimates of several different discount functions into a single measure of the discount rate, in effect treating all estimates as if they were generated by an exponential model. This allows them to quantify the *magnitude* of discounting, but not its *source*. By conducting our meta-analysis within the framework of the quasi-hyperbolic model, we are able to distinguish the two sources recognised within that model, namely present bias and long-run impatience. We find support for time inconsistency as viewed through the lens of this model, and especially so for non-monetary rewards. Nonetheless, we reiterate that to evaluate the β - δ model against

¹¹For the avoidance of doubt, we clarify that this usage of the symbol τ by [Bleichrodt et al. \(2022\)](#) is unrelated to our usage of the same notation to denote the heterogeneity parameter of the random-effects meta-analysis model in Section 2.3.1.

other alternative formulations calls for a richer design than most of the ones reviewed here. For our own first steps in this direction, we refer the interested reader to [Cheung et al. \(2024\)](#).

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.
- Ainslie, G., Haendel, V., 1983. The motives of the will, in: Gottheil, E., Druley, K.A., Skoloda, T.E., Waxman, H.M. (Eds.), *Etiologic Aspects of Alcohol and Drug Abuse*. Charles C. Thomas, Springfield, IL, pp. 119–140.
- 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., Kuhn, M.A., Sprenger, C., 2015. Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior and Organization* 116, 451–464.
- 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.
- Bleichrodt, H., Potter van Loon, R.J., Prelec, D., 2022. Beta-Delta or Delta-Tau? A reformulation of quasi-hyperbolic discounting. *Management Science* 68, 6326–6335.
- Bleichrodt, H., Rohde, K.I., Wakker, P.P., 2009. Non-hyperbolic time inconsistency. *Games and Economic Behavior* 66, 27–38.
- Bom, P.R., Rachinger, H., 2019. A kinked meta-regression model for publication bias correction. *Research Synthesis Methods* 10, 497–514.
- 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.
- Brown, A.L., Vieider, F.M., Imai, T., Camerer, C.F., 2024. Meta-analysis of empirical estimates of loss aversion. *Journal of Economic Literature* 62, 485–516.
- 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., MacGibbon, K., Milin-Byrne, A., Tymula, A., 2024. Quasi-exponential discounting. <https://ssrn.com/abstract=4910013>.
- Cheung, S.L., Tymula, A., Wang, X., 2022. Present bias for monetary and dietary rewards. *Experimental Economics* 25, 1202–1233.

- Cohen, J., Ericson, K.M., Laibson, D., White, J.M., 2020. Measuring time preferences. *Journal of Economic Literature* 58, 299–347.
- Coller, M., Williams, M.B., 1999. Eliciting individual discount rates. *Experimental Economics* 2, 107–127.
- DellaVigna, S., Malmendier, U., 2006. Paying not to go to the gym. *American Economic Review* 96, 694–719.
- DerSimonian, R., Laird, N., 1986. Meta-analysis in clinical trials. *Controlled Clinical Trials* 7, 177–188.
- Ebert, J.E., Prelec, D., 2007. The fragility of time: Time-insensitivity and valuation of the near and far future. *Management Science* 53, 1423–1438.
- Egger, M., Davey Smith, G., Schneider, M., Minder, C., 1997. Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal* 315, 629–634.
- Eicher, T.S., Papageorgiou, C., Raftery, A.E., 2011. Default priors and predictive performance in Bayesian model averaging, with application to growth determinants. *Journal of Applied Econometrics* 26, 30–55.
- Estle, S.J., Green, L., Myerson, J., Holt, D.D., 2007. Discounting of monetary and directly consumable rewards. *Psychological Science* 18, 58–63.
- Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* 60, 351–401.
- George, E.I., 2010. Dilution priors: Compensating for model space redundancy, in: *Borrowing Strength: Theory Powering Applications – A Festschrift for Lawrence D. Brown*. Institute of Mathematical Statistics, pp. 158–165.
- Goldstein, H., Yang, M., Omar, R., Turner, R., Thompson, S., 2000. Meta-analysis using multilevel models with an application to the study of class size effects. *Journal of the Royal Statistical Society Series C: Applied Statistics* 49, 399–412.
- Halevy, Y., 2015. Time consistency: Stationarity and time invariance. *Econometrica* 83, 335–352.
- Hedges, L., Olkin, I., 1985. *Statistical Methods for Meta-Analysis*. Academic Press.
- Hedges, L.V., 1983. A random effects model for effect sizes. *Psychological Bulletin* 93, 388–395.
- Higgins, J., Thompson, S., 2002. Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine* 21, 1539–1558.
- Imai, T., Rutter, T.A., Camerer, C.F., 2021. Meta-analysis of present-bias estimation using convex time budgets. *Economic Journal* 131, 1788–1814.
- Ioannidis, J.P., Stanley, T., Doucouliagos, H., 2017. The power of bias in economics research. *Economic Journal* 127, F236–F265.
- Irsova, Z., Bom, P.R., Havranek, T., Rachinger, H., 2023. Spurious precision in meta-analysis. Discussion Paper 17927, Centre for Economic Policy Research.
- Irsova, Z., Doucouliagos, H., Havranek, T., Stanley, T.D., 2024. Meta-analysis of social science research: A practitioner's guide. *Journal of Economic Surveys* 38, 1547–1566.

- Jackson, D., White, I.R., Riley, R.D., 2020. Multivariate meta-analysis, in: Schmid, C.H., Stijnen, T., White, I.R. (Eds.), *Handbook of Meta-Analysis*. Chapman and Hall, pp. 163–186.
- Kaur, S., Kremer, M., Mullainathan, S., 2015. Self-control at work. *Journal of Political Economy* 123, 1227–1277.
- Laibson, D., 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112, 443–478.
- Laibson, D., Repetto, A., Tobacman, J., 1998. Self-control and saving for retirement. *Brookings Papers on Economic Activity*, 91–196.
- Light, R.J., Pillemer, D.B., 1984. *Summing Up: The Science of Reviewing Research*, Harvard University Press, Cambridge, MA.
- Lipman, S.A., Attema, A.E., 2024. A systematic review of unique methods for measuring discount rates. *Journal of Risk and Uncertainty* 69, 145–189.
- Loewenstein, G.F., Prelec, D., 1992. Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics* 107, 573–597.
- Matousek, J., Havranek, T., Irsova, Z., 2022. Individual discount rates: A meta-analysis of experimental evidence. *Experimental Economics* 25, 318–358.
- Mazur, J., 1987. An adjusting procedure for studying delayed reinforcement, in: Commons, M., Mazur, J., Nevin, J., Rachlin, H. (Eds.), *The effect of delay and of intervening events on reinforcement value*, 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.
- Rachlin, H., 2006. Notes on discounting. *Journal of the Experimental Analysis of Behavior* 85, 425–435.
- Raftery, A.E., Madigan, D., Hoeting, J.A., 1997. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association* 92, 179–191.
- Raudenbush, S., 2009. Analyzing effect sizes: Random-effects models, in: Cooper, H., Hedges, L., Valentine, J. (Eds.), *The Handbook of Research Synthesis and Meta-Analysis*. 2nd ed.. Russell Sage Foundation, pp. 295–316.
- Read, D., 2001. Is time discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty* 23, 5–32.
- Reuben, E., Sapienza, P., Zingales, L., 2010. Time discounting for primary and monetary rewards. *Economics Letters* 106, 125–127.
- Samuelson, P.A., 1937. A note on measurement of utility. *Review of Economic Studies* 4, 155–161.
- Scholten, M., Read, D., 2006. Discounting by intervals: A generalized model of intertemporal choice. *Management Science* 52, 1424–1436.
- Scholten, M., Read, D., 2010. The psychology of intertemporal tradeoffs. *Psychological Review* 117, 925–944.

- Stanley, T., Doucouliagos, H., 2014. Meta-regression approximations to reduce publication selection bias. *Research Synthesis Methods* 5, 60–78.
- Steel, M.F., 2020. Model averaging and its use in economics. *Journal of Economic Literature* 58, 644–719.
- Strotz, R., 1955. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23, 165–180.
- Thaler, R., 1981. Some empirical evidence on dynamic inconsistency. *Economics Letters* 8, 201–207.
- Thompson, S.G., Turner, R.M., Warn, D.E., 2001. Multilevel models for meta-analysis, and their application to absolute risk differences. *Statistical Methods in Medical Research* 10, 375–392.
- Ubfal, D., 2016. How general are time preferences? Eliciting good-specific discount rates. *Journal of Development Economics* 118, 150–170.