



City Research Online

City St George's, University of London

Citation: Cerrone, C., Chakraborty, A., Kim, H. J. & Lades, L. K. (2025). Estimating present bias and sophistication over effort and money. *Experimental Economics*, 28(6), pp. 1191-1213. doi: 10.1017/eec.2026.10053

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/37956/>

Link to published version: <https://doi.org/10.1017/eec.2026.10053>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

ORIGINAL PAPER

Estimating present bias and sophistication over effort and money

Claudia Cerrone¹ , Anujit Chakraborty² , Hyok Jung Kim³  and Leonhard K. Lades⁴ 

¹Department of Economics, City St George's, University of London

²Department of Economics, University of California, Davis

³Department of Economics, Korea Institute for International Economic Policy

⁴Economics and Behavioural Science at the Stirling Business School, University of Stirling

Corresponding author: Claudia Cerrone; Email: claudia.cerrone@city.ac.uk

(Received 6 October 2024 UTC; revised 24 September 2025 UTC; accepted 2 May 2026 UTC)

Abstract

We design and conduct a real-effort experiment to jointly estimate present bias and sophistication across effort and monetary domains. Unlike prior work (e.g., Augenblick and Rabin, 2019; Fedyk, 2024), we do not assume that these parameters are identical across domains. We explain and empirically demonstrate how assuming identical sophistication across money and effort domains can bias the estimates of key parameters. In our online experiment, participants chose to (predicted to) complete 14% (10%) fewer tasks on the same day than on a future day, leading to an estimated present bias (β_e) over effort of 0.70–0.79, and an estimated sophistication ($\hat{\beta}_e$) of 0.80–0.88. For money, aggregate present bias (β_m) is near zero, but there is substantial heterogeneity, with roughly equal numbers of participants exhibiting present bias and future bias. At the individual level, roughly three quarters of all participants correctly anticipate the direction of their bias in both domains, even if not its full magnitude.

Keywords: Experiment; present bias; sophistication; time preferences

JEL Codes: C91; D81; D90

1. Introduction

Present bias – the desire for immediate gratification – often derails long-term goals such as saving for retirement, starting a diet, or finishing a referee report. Individuals who are present-biased and *sophisticated*, i.e., aware of their own present bias, can avoid making present-biased decisions by limiting the options they will have in the future (O'Donoghue & Rabin, 1999). For example, they might self-impose stringent work deadlines (Ariely & Wertenbroch, 2002), invest in illiquid savings products (Ashraf et al., 2006), sign up for effort-incentivizing contracts (Kaur et al., 2015), or use commitment devices that restrict their future choices (Schilbach, 2019).¹ Predicting the dynamic decisions of present-biased individuals, therefore, requires the joint measurement of present bias and sophistication. Yet many empirical studies measure present bias while ignoring sophistication.

The only existing experimental design that directly estimates sophistication, introduced in Augenblick and Rabin (2019) and extended by Fedyk (2024), represents an important contribution. In

¹Toussaert (2018) shows that experimental participants who anticipate self-control costs as in Gul and Pesendorfer (2001) may commit to a smaller choice set even when they expect to resist temptation.

these designs, participants make real-effort choices at varying wages and predict their future choices; sophistication can be measured by the gap between predictions and actual behaviour.

The identification strategy in these studies, however, relies on the assumption that present bias and sophistication are *identical* across effort and money domains. This restriction is substantive as prior work documents systematic domain differences. Empirically, present bias has been robustly documented in the effort domain but is often small or statistically insignificant in the money domain (Augenblick et al., 2015; Imai et al., 2021). Conceptually, money is fungible and can be stored or reallocated, whereas effort is non-fungible, immediately experienced, and cannot be deferred or transferred. These differences imply that behaviour associated with exerting effort may be qualitatively distinct from behaviour associated with receiving monetary payments. As we explain in Online Appendix C, imposing a common parameter across domains therefore conflates present bias and sophistication over money and effort, yielding estimates that may fail to capture behaviour in either domain.

Our experiment, conducted online with Prolific participants over 3 days, builds directly on the approach of Augenblick and Rabin (2019) and Fedyk (2024) but introduces a design feature that enables separate identification of present bias and sophistication over money and effort. The key novelty is that all payments are made on Day 3 and this timing is made fully salient to participants. This creates two independent sources of variation. Firstly, we observe choices and predictions when effort costs are immediate (decisions for tasks to be completed on the same day) and when they are delayed (decisions for tasks to be completed on a future day). Comparing choices (and predictions) under immediate versus delayed effort identifies present bias (and sophistication) over effort. Secondly, we observe choices and predictions when monetary benefits are effectively immediate (decisions made on Day 3, paid on Day 3) and when they are delayed (decisions made on Days 1 or 2, paid on Day 3). Comparing choices (and predictions) under immediate versus delayed payment identifies present bias (and sophistication) over money. Thus, our double experimental variation allows us to separately identify present bias and sophistication in the effort and money domains. This contrasts with the designs in Augenblick and Rabin (2019); Fedyk (2024), where all payments are in the future, and hence identifying crucial model parameters requires assuming identical present bias and sophistication across domains.

We begin with a reduced-form analysis, which provides significant evidence for present bias in the effort domain, but little to no evidence for present bias over money. Controlling for the distance between payment and effort dates, participants allocate on average 4.17 (14%) more tasks to future workdays than to the present, but only 0.5 fewer tasks when payment occurs in the future rather than the present – only the former effect is statistically significant. When predicting their own future decisions, participants believe they will allocate 3.09 (rather than 4.17) more tasks to future workdays, providing statistically significant evidence of partial sophistication over effort.

For our structural analysis, we use the quasi-hyperbolic discounting model (Phelps & Pollak, 1968; Laibson, 1997) due to its tractability.² We find significant evidence of sophistication over effort ($\hat{\beta}_e$): across specifications, the point estimate of $\hat{\beta}_e$ ranges between 0.80 and 0.88, while the point estimate of present bias over effort (β_e) lies between 0.70 and 0.79. Our aggregate estimate of β_m ranges between 0.95 and 1.09 and is not statistically different from 1, though substantial individual-level heterogeneity exists, with roughly equal numbers of participants exhibiting present bias and future bias over money. Overall, roughly three quarters of all participants anticipate the direction of their bias in both domains, even if not its full magnitude.

To assess the cost of imposing the common-parameter assumption, we re-estimate our data under Augenblick and Rabin (2019)'s original restriction. This counterfactual exercise confirms that the

²See Chakraborty (2021) for a detailed discussion of the many ways present-biased preferences can be modelled.

assumption biases estimates of key preference parameters: it forces β and $\hat{\beta}$ mutually closer together than in our preferred specifications, thereby overstating relative sophistication.

Besides serving as a proof of concept for the identification strategy, our 3-day experimental design is more cost-effective, less vulnerable to attrition, and easier to replicate than longer studies. It can moreover be extended to longer horizons or field settings, making it a useful tool for applied researchers who wish to use present bias and sophistication as explanatory variables in domains such as savings behaviour, illiquid investments, gym memberships, and workplace goal-setting.

The remainder of the paper is structured as follows. [Section 2](#) reviews the related literature and situates our contribution within this literature. [Section 3](#) describes the experimental design. [Section 4](#) introduces the model and explains the identification strategy. [Section 5](#) presents our results. [Section 6](#) compares our results with those in closely related previous studies. [Section 7](#) concludes.

2. Literature review

2.1. Measuring sophistication

The literature measuring sophistication about time preferences can be organized into three broad strands. The first strand infers sophistication from individuals' preferences over commitment devices or restricted choice sets. Ahn et al. (2019) provide a behavioural characterization of naïvete based on ex-ante choice over menus and ex-post choice from menus: an individual is sophisticated if she is indifferent ex ante between retaining the option to choose from a menu ex post or committing to her actual distribution of choices from that menu. Other work in this strand treats demand for commitment as a 'smoking gun' for sophistication (e.g., Ericson & Laibson, 2019; Sadoff & Samek, 2019; Sprenger, 2015). But commitment demand is only an outcome of sophistication, not a direct measure. Moreover, Carrera et al. (2022) show that experimentally increasing sophistication can actually reduce demand for commitment, contradicting the assumption that commitment demand cleanly reveals sophistication (see also Westphal, 2024).

The second strand measures sophistication more directly by eliciting beliefs about future behaviour and comparing them to realized behaviour. Augenblick and Rabin (2019) and Fedyk (2024) use incentivized real-effort experiments in which participants decide how much to work at different times under varying wages; sophistication is measured by comparing actual work decisions on future dates with participants' earlier predictions of those same decisions. Both studies find that participants are naïve about their own present bias: they fail to fully anticipate it, leading them to work less than they had predicted. Fedyk (2024) further shows that participants are more sophisticated about others' present bias than about their own, a phenomenon they call 'asymmetric naïvete'.

Survey-based approaches also belong to this strand. For example, Cobb-Clark et al. (2024) classify individuals according to the consistency between their ideal, expected, and actual body weight 1 year later. While intuitive, such proxies can misclassify behaviour. For instance, someone who predicts they will not reach their ideal weight would be labelled sophisticated, even if they are simply unmotivated by the goal. More fundamentally, such survey-based classifications do not yield structural parameters such as β and $\hat{\beta}$. Sial et al. (2023) link forecast errors to memory biases, finding that individuals who over-report past attendance are less aware of their time inconsistency. The present paper belongs to this second strand, using incentivized measures of precise behaviours and predictions thereof, while extending prior work by allowing present bias and sophistication to differ across effort and money domains.

The third strand relies on proxies or unincentivized measures of sophistication, such as cognitive ability or related psychological traits. For example, Zhang and Greiner (2021) link cognitive ability to awareness of present bias.

2.2. Identification of time preference parameters

A distinct but related body of work addresses identification challenges in estimating time-preference parameters from intertemporal choice data. Strack and Taubinsky (2026) highlight important challenges in ‘revision design’ experiments, where participants make intertemporal trade-offs between two periods, deciding first in advance and then again when the time comes. They show that with only ordinal choice data, the same observed behaviour can be rationalized by multiple assumptions about how information about future preferences or constraints is revealed to the decision-maker. They note that when trade-offs are measured in cardinal terms, such as monetary amounts, identification becomes possible. Our design exploits this insight by eliciting choices over monetary amounts across time, enabling us to more precisely recover present bias and sophistication by leveraging variation in participants’ willingness to pay.

2.3. Domain-sensitive time preference parameters

Previous literature shows that time preferences are domain-sensitive. People tend to be more present-biased over primary rewards such as food, juice, or work, while present bias often appears weaker – or even absent – for monetary rewards. For example, Cheung et al. (2022) show that present bias differs across money, healthy foods, and unhealthy foods.

Of particular relevance to our paper, this domain-sensitivity extends to the money-effort distinction, with evidence suggesting that present bias differs systematically between the two. For example, Augenblick et al. (2015) find no present bias over money but significant present bias over effort. In a meta-analysis about studies using the Convex Time Budget method, Imai et al. (2021) show that studies using monetary rewards often show statistically insignificant present bias on average, while effort-based studies tend to report a smaller and significant present bias parameter. Identifying sophistication in Augenblick and Rabin (2019) and Fedyk (2024) requires *assuming* equal present bias and sophistication over money and effort, thereby yielding estimates that may fail to capture behaviour in either domain. This is what distinguishes the current paper: our experimental design identifies sophistication and effort, without assuming equality across money and effort domains.

3. Experimental design

3.1. Sessions

Fig. 1 illustrates the sequence of events across the 3 consecutive days of a session. On Day 1’s consent form, participants were informed that the study involved real-effort tasks over 3 days. Participants then read the instructions and completed a comprehension quiz assessing their understanding of the study’s structure – namely, its 3-day duration, the £5 bonus for completing all days, and the fact that all payments (for tasks and study completion) would be made on Day 3. Before making any decisions, they were required to complete training tasks to familiarize themselves with the task format and its difficulty. The information about study and payment structures was reiterated on Days 2 and 3 as a reminder, then they made all decisions and predictions, and then they completed any effort tasks scheduled for that day.

3.2. Decisions and predictions

Participants were asked to choose and predict the number of time-dated real-effort tasks that they wanted to complete for different wages on the current day and on future days.

Over the 3 days, participants saw 9 decision screens similar to the one shown in Fig. 2. On each of the 9 screens, participants saw 5 different wages or piece rates (1p, 3p, 5p, 7p, and 9p) and had to choose the number of time-dated *jobs*, between 3 and 13, that they wanted to complete at that wage and on that date. Six of the 9 screens were choice screens, on which participants made their




Fig. 1. Session timeline: the timeline of events within each of the 3 days

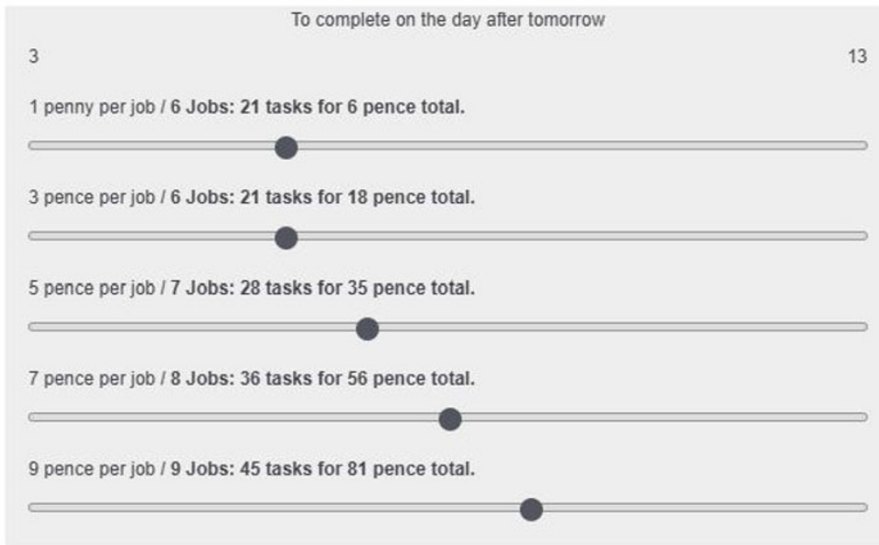


Fig. 2. Decision screen for choosing jobs $J \in \{1, 2, \dots, 13\}$ at 5 wages. At each wage w , for the selected J , the screen dynamically calculated and displayed the number of tasks $J(J+1)/2$ and the total corresponding payment wJ . For example, in the figure above, the participant selected $J = 3$ at $w = 1$, and hence the screen said ‘3 Jobs: 6 tasks for 3 pence total’

own decisions, while the remaining 3 were prediction screens, on which they predicted the number of jobs their future selves would choose to complete. This means that participants made $5 \times 9 = 45$ decisions in total.

Jobs translated into real-effort tasks at a non-linear rate: J jobs required completing $J(J+1)/2$ real-effort transcription tasks (just *tasks*, henceforth). For example, as seen on Fig. 2, $J = 3$ jobs required completing $3(3+1)/2 = 6$ tasks, $J = 5$ jobs required completing $5(5+1)/2 = 15$ tasks, and $J = 13$ jobs required completing $13(13+1)/2 = 91$ tasks.

As illustrated by Fig. 3, a (transcription) task consisted of typing a sequence of characters into an empty box. If the characters were not typed correctly, an error message appeared and the participant had to retype the characters again. Our transcription tasks resemble those used in Augenblick and Rabin (2019).³

3.2.1. Increasing marginal cost of jobs

Participants were paid the wage or piece rate based on the number of jobs completed. Thus, even though every job offered the same marginal benefit of piece rate p , the J -th job required completing

³In Fedyk (2024), participants were instead asked to press a key every time an asterisk appeared in a random sequence of characters.



Fig. 3. Example of a transcription task. Completing J jobs meant finishing $J(J + 1)/2$ transcription tasks for a total payment of Jp

Table 1. Decisions and predictions by date

		Work date (τ)		
		1	2	3
Decision date (t)	1	$J_{11}(w)$	$J_{12}(w), J_{22}^p(w)$	$J_{13}(w), J_{23}^p(w)$
	2		$J_{22}(w)$	$J_{23}(w), J_{33}^p(w)$
	3			$J_{33}(w)$

$J_{t\tau}(w)$ denotes the decision made on date t about the number of jobs to do on date τ . $\widehat{J}_{t\tau}(w)$ denotes the prediction about $J_{t\tau}(w)$. $w \in \{1, 3, 5, 7, 9\}$ denotes the wage rate.

$J(J + 1)/2 - (J - 1)J/2 = J$ additional real-effort tasks on the margin. This means that participants have to work more and longer to complete each subsequent job.⁴ This design feature helps aligning the job choice with the canonical assumption of the cost being convex and the net utility being concave in jobs.

Theoretically, this feature should also decrease the frequency of corner solutions, i.e., fewer participants should choose either the minimum or the maximum amount of tasks to complete. The reduction of corner choices makes it easier to estimate structural parameters based on the assumption of internal optima.⁵ The mapping from jobs to tasks was explained in advance, as well as transparently presented to the participants on every decision screen (see Fig. 2).

3.2.2. All choices and predictions

Table 1 summarizes all the decisions made by participants by date. $J_{t\tau}(w)$ denotes the decisions made on day t about the number of jobs to complete on day τ at the wage w . On Day 1, participants decided the number of jobs to complete on Days 1, 2, and 3, denoted by $J_{11}(w)$, $J_{12}(w)$, and $J_{13}(w)$, respectively. On Day 2, they decided the number of jobs to complete on Days 2 and 3, i.e., $J_{22}(w)$ and $J_{23}(w)$, respectively. On Day 3, they decided how many jobs to complete on Day 3, denoted by $J_{33}(w)$.

On each day τ , one decision was randomly chosen from the pool of decisions made for that day, and participants had to complete the corresponding number of jobs for the corresponding wage. As an example, on Day 3, one decision was randomly drawn from a pool of 15 decisions: decisions J_{13} , J_{23} , and J_{33} , each of them made for 5 different wages.

After making their decisions, participants completed the tasks corresponding to the randomly selected decision for that day. Only after completing these tasks did they proceed to the prediction stage. This temporal separation ensures that decisions and predictions are meaningfully separated in time.

$J_{t\tau}^p(w)$ denotes the prediction of the decision $J_{t\tau}(w)$. Participants faced 3 prediction screens. On Day 1, they made predictions $J_{22}^p(w)$ and $J_{23}^p(w)$ about $J_{22}(w)$ and $J_{23}(w)$, respectively. On Day 2, they predicted $J_{33}(w)$.

⁴For comparison, in Augenblick and Rabin (2019) and Chakraborty and Fenig (2025), the total earnings increase proportionally with effort.
⁵See Le Yaouanq and Schwardmann (2022) and Gerhardt and Suchy (2024) for similar approaches.

3.3. Payments

All payments were made within the evening of the third day after the survey was over, as promised in the instructions. To make sure that participants understood all the sources of payment, we quizzed them about the experiment's payment structure, including one question about the timing of the payments.⁶

There were four sources of payment for each participant. Firstly, participants received a participation fee for each day they participated in the study (£2, £1, and £1, respectively). The participation fee was higher on the first day than on the other 2 days as participants spent more time reading instructions, completing mandatory tasks for training, and making more decisions. Secondly, participants received a bonus of £5 if they completed all 3 days of the experiment. This completion fee was relatively high to minimize attrition. Thirdly, on each day, participants completed tasks and were paid according to one randomly chosen decision out of the decisions they made for that day. To calculate this payment, the randomly drawn wage rate was multiplied by the number of jobs chosen and completed for that decision. On each day, the payment could range from 3p ($1p \times 3$ jobs) to 117p ($9p \times 13$ jobs). On the 3 survey days, participants earned 35p, 35p, and 32p on average for completing their chosen number of transcription tasks, respectively.

Finally, participants could receive a fourth payment of 20p if one of their predictions (randomly chosen from their 15 predictions) differed by no more than 2 jobs from their actual decision. Thus, on average, every accurate prediction task only paid 1.33p. Around 87% of all predictions were accurate within 2 jobs, and hence qualified for the 20p reward if chosen for payment. Unlike Augenblick and Rabin (2019), who reminded participants of their prior predictions before final decisions, we follow Fedyk (2024) in omitting such reminders to avoid priming participants to behave consistently with earlier responses.

Setting the prediction reward involved a trade-off. Too large an accuracy bonus would allow sufficiently sophisticated participants to use the prediction tasks as a commitment device, since the prospect of matching a previously made prediction could incentivize their future selves to work harder than wage incentives alone would warrant. Too small a bonus, on the other hand, might lead participants to provide careless predictions.⁷ We therefore set a very small accuracy bonus as a compromise between these two concerns.

3.4. Participants and selection

All experimental sessions were run online in November 2021 and May 2022. With the goal of having a sample of around 100 participants post-attrition, we recruited 138 participants through Prolific. A total of 115 participants completed the study, although 7 had some missing data. In November, the median length of the experiment was 30 minutes on the first day, 15 minutes on the second day, and 9 minutes on the third day. In May, the median length was 27, 12, and 12 minutes, respectively. A total of 108 participants completed all 3 days of the experiment. This study was not preregistered, as the data were collected prior to the widespread adoption of preregistration practices in the field.

Out of the 138 participants who completed the first day of the study, 108 responded to all questions during the 3 days. Of these, 7 participants did not show any variation across any of the 9 screens, leaving 101 participants for meaningful regression analyses. To winsorize this sample, we further excluded the participants who ranked in the top 10% of 'most frequent violators' in each of the following categories: wage monotonicity, time-monotonicity with respect to effort date, and time-monotonicity with respect to pay date. Our final dataset therefore consists of 80 participants (out of

⁶ After verifying a new participant's credentials and the first four cash outs, Prolific allows participants to instantly cash out their payments from all future studies. There were no complaints about the payments not going through on time.

⁷ A strong argument in favour of providing small incentives for predictions comes from Augenblick and Rabin (2019), who randomly assigned accuracy bonuses between 0.25 and 8.25 to identify if higher accuracy payments affected predictions or future decisions, and found no such evidence.

the 101 regression-eligible). For consistency, and if not otherwise stated, all the tables or graphs that follow use this sample of 80 participants.

The three monotonicity conditions used to define the sample restrictions – wage monotonicity, time-monotonicity with respect to effort date, and time-monotonicity with respect to pay date – are discussed in detail, along with the corresponding statistics, in Online Appendix A.

4. Model and identification strategy

To motivate our analysis of the data, we start with a general model of behaviour where participants (implicitly) trade off the benefits of completing jobs (i.e., the money, Jw) with the costs (i.e., the effort that must be exerted). We assume that on decision day t_d , participants choose their optimal effort J^* for day $t_e \geq t_d$ on the experimental task with piece-rate w based on the following utility optimization exercise:

$$J^* = \arg \max_J \psi \beta_m^{\mathbb{1}(3 > t_d)} \delta^{3-t_d} Jw - \beta_e^{\mathbb{1}(t_e > t_d)} \delta^{t_e-t_d} c(J(J+1)/2) \quad (1)$$

where $c(\cdot)$ is the cost from the $J(J+1)/2$ transcription tasks that need to be completed for J jobs, as previously described in Section 3.2. The indicator $\mathbb{1}(3 > t_d)$ denotes whether the payment day (Day 3) is later than the decision day t_d . Similarly, $\mathbb{1}(t_e > t_d)$ indicates whether the decision is being made about working in the future, and $(t_e - t_d)$ measures the distance between the effort date and the decision date. The scale parameter ψ normalizes the benefit of monetary gains and the cost from transcription tasks at the same unit.⁸

When asked to predict the decision J^* at some prediction period $t_p < t_d$, participants report:

$$\hat{J} = \arg \max_J \psi \hat{\beta}_m^{\mathbb{1}(3 > t_d)} \delta^{3-t_d} Jw - \hat{\beta}_e^{\mathbb{1}(t_e > t_d)} \delta^{t_e-t_d} c(J(J+1)/2) \quad (2)$$

To get to Equation ((2)) from ((1)), one simply replaces β_m by $\hat{\beta}_m$ and β_e by $\hat{\beta}_e$. In Equation ((2)), the first term is the perceived discounted benefit of completing more jobs and the second term is its perceived discounted cost.

4.1. Time-consistent vs. sophisticated vs. naïve

Following O'Donoghue and Rabin (1999), we distinguish between three mutually exclusive types of individuals. When $\beta_m = \beta_e = 1$, the individual has exponential discounting preferences. In this case, her relative valuation of earlier versus later utility remains constant regardless of when she is asked to make that trade-off. Since there is no conflict between her preferences at different points in time, such a person is said to be *time-consistent*, and her plans/predictions automatically align with her future actions.

When either $\beta_m < 1$ or $\beta_e < 1$, the individual exhibits present bias. When considering trade-offs of utility between two moments, present-biased individuals give stronger relative weight to the earlier moment if it is also in the present. This bias creates the possibility of a divergence between plans/predictions and future actions, depending on whether the individual anticipates her present bias.

O'Donoghue and Rabin (1999) classify individuals as sophisticated or naïve only when they are present-biased (i.e., $\beta_m < 1$ or $\beta_e < 1$). If the individual correctly anticipates her present bias, she is *sophisticated*; otherwise, she is *naïve*.

More generally, we interpret sophistication as the accuracy of individuals' beliefs about their own present-bias parameters, $\hat{\beta}_m$ and $\hat{\beta}_e$, which can be defined irrespective of whether $\beta < 1$ or $\beta \geq 1$. A fully sophisticated individual satisfies $\hat{\beta}_e = \beta_e$ in the effort domain (or $\hat{\beta}_m = \beta_m$ in the money domain), while deviations between $\hat{\beta}$ and β reflect the degree of naïveté. For completeness, our graphs

⁸For simplicity, we assume that δ is equal for money and effort.

and tables report both present bias (β) and perceived present bias ($\hat{\beta}$) for all participants, including those with $\beta \geq 1$.

4.2. Identification

Overall, as long as $c(\cdot)$ is convex in J , the problem in Equation ((1)) or ((2)) is strictly concave and has a unique solution. In particular, for our parametric estimation exercise, we will assume that $c(x) = \frac{x^{\gamma+1}}{\gamma+1}$ and substitute it in the first-order condition derived from (2):

$$\left(\psi \frac{\hat{\beta}_m^{1(3>t_d)} w \delta^{3-t_e}}{\hat{\beta}_e^{1(t_e>t_d)}} \right) - \frac{2\hat{J}+1}{2} c'(\hat{J}(\hat{J}+1)/2) = 0. \quad (3)$$

$$\Leftrightarrow \left(\psi \frac{\hat{\beta}_m^{1(3>t_d)} w \delta^{3-t_e}}{\hat{\beta}_e^{1(t_e>t_d)}} \right) - \left(\frac{\hat{J}(\hat{J}+1)}{2} \right)^{\gamma} \left(\frac{2\hat{J}+1}{2} \right) = 0. \quad (4)$$

Thus, the prediction $\hat{J}$ directly reflects the individual's beliefs about her own future preferences, as captured by $\hat{\beta}_m$ and $\hat{\beta}_e$. This identification strategy relies on other components – such as the marginal cost function $c'(\cdot)$ – being known. This dependency underscores the importance of the training tasks administered before any decisions were made on Day 1 (see Fig. 1), as they were designed to increase participants' familiarity with both the overall cost and the marginal cost of performing the tasks. As before, one can move between the optimality conditions for J^* and $\hat{J}$ by interchanging β_m with $\hat{\beta}_m$, and β_e with $\hat{\beta}_e$.

4.3. Our approach

To jointly estimate β and $\hat{\beta}$ parameters *separately* for effort and money, our experimental design allows us to independently observe:

- **Variation 1:** Choices and predictions when the effort costs are in the future (tasks to be completed on later days) *versus* when the effort costs are immediate (tasks to be completed the same day).
- **Variation 2:** Choices and predictions when the benefits (payments) are in the future (decisions made on Days 1 and 2, paid on Day 3) *versus* when the benefits are immediate (decisions made on Day 3, paid the same day).

While other studies such as Augenblick and Rabin (2019) and Fedyk (2024) also use Variation 1, they lack Variation 2, as their payment date is always strictly after every effort date. Thus, β_m and $\hat{\beta}_m$ can never be identified separately from β_e and $\hat{\beta}_e$, respectively.

Equipped with Variations 1 and 2, and assuming i.i.d. normally distributed errors with equal variance across choice and predictions, we estimate the parameter vector of the model, i.e., $\theta = (\beta_e, \hat{\beta}_e, \beta_m, \hat{\beta}_m, \delta, \psi, \gamma)$, using the maximum likelihood estimator (MLE) over both choice and prediction data. $\hat{\theta}_{MLE}$, the maximum likelihood estimate of θ , is defined as

$$\hat{\theta}_{MLE} = \arg \max_{\theta} \sum_{f=1}^F \ln \mathcal{L}_f(\theta), \quad \text{where}$$

$$\ln \mathcal{L}_f(\theta) = \begin{cases} -\ln \sigma - 0.5 \ln 2\pi - 0.5 \left(\frac{J_f^*(\theta) - J_f}{\sigma} \right)^2 & \text{For choice data} \\ -\ln \sigma - 0.5 \ln 2\pi - 0.5 \left(\frac{J_f(\theta) - J_f}{\sigma} \right)^2 & \text{For prediction data} \end{cases}$$

where f is an index for choice and prediction data points, combined across all individuals and 3 days for aggregate analysis, and fixing an individual over 3 days for individual-level analysis. For any value of θ , $J_f^*(\theta)$, $\widehat{J_f}(\theta)$ are calculated using Equation (4). MLE finds the set of parameters θ that maximizes the joint likelihood of observing the complete data.

5. Results

For the data to be consistent with any rational model of choice, the number of jobs chosen (or predicted to be chosen) should weakly increase with the wage, and should also weakly increase when the work is scheduled for a later day. Online Appendix A shows that these monotonicity conditions (Chakraborty et al., 2017) hold for the vast majority of our data. Moreover, we also discuss the frequency of corner choices in the data.

5.1. Reduced-form analysis

Following Augenblick and Rabin (2019), we use the following regression specification at both the aggregate level (Section 5.1.1) and the individual level (Section 5.1.2) to find the marginal effect of present bias over money and effort on effort choices, without imposing any parametric assumptions on the felicity function or discounting. Our regression specification is:

$$e_{i,j} = \phi_0 + \phi_e^{\text{future}} 1_{i,j}^{\text{future } e} + \phi_e^{\text{predict}} 1_{i,j}^{\text{predict } e} + \phi_m^{\text{future}} 1_{i,j}^{\text{future } m} + \phi_m^{\text{predict}} 1_{i,j}^{\text{predict } m} + \phi_\gamma w_{i,j} + \phi_\delta (3 - t_e) + \mu_i + \varepsilon_{i,j}. \quad (5)$$

where $e_{i,j}$ is the number of jobs (or tasks) chosen by individual i at observation j . Observation j can be either an effort decision or an effort prediction. The regression controls for standard impatience or discounting through $(3 - t_e)$, the relative distance between the payment date and the effort date.⁹ To allow for different base levels of efforts chosen by different individuals, we include participant-level fixed effect terms μ_i as dummy variables. We also control for the wage rate $w_{i,j}$: $\phi_\gamma > 0$ captures the wage monotonicity. Conditional on observation j being an effort decision, $1_{i,j}^{\text{future } e}$ equals one if j involves future effort, and $1_{i,j}^{\text{future } m}$ equals one if the payment is made in the future. $1_{i,j}^{\text{predict } e}$ and $1_{i,j}^{\text{predict } m}$ are the corresponding indicators when observation j is an effort prediction instead.

For a fixed relative distance $(3 - t_e)$ between the payment date and the effort date, ϕ_e^{future} measures the additional number of jobs that participants choose to complete in the future, relative to the present. A positive value of ϕ_e^{future} implies participants choose more effort if it needs to be exerted in the future than if it needs to be exerted immediately, indicating a present bias over effort. The higher is ϕ_e^{future} , the higher the implied present bias over effort (i.e., the lower is β_e). Thus, $-\phi_e^{\text{future}}$ provides a reduced-form measure that maps to β_e , with lower values of either corresponding to more present-biased individuals.

Similarly, $-\phi_e^{\text{predict}}$ provides a reduced-form measure of *perceived* present bias over effort. Sophistication is reflected in how closely perceived present bias aligns with actual present bias.

In the money domain, ϕ_m^{future} measures how effort choices respond to the timing of payments. A negative value of ϕ_m^{future} indicates that participants choose fewer jobs when payments are immediate rather than delayed, consistent with present bias over money (i.e., $\beta_m < 1$). Thus, ϕ_m^{future} provides a reduced-form measure of actual present bias over money, with lower values of either corresponding to more present-biased individuals. Similarly, ϕ_m^{predict} captures perceived present bias over money, while the difference between ϕ_m^{future} and ϕ_m^{predict} reflects the degree of relative sophistication over money (β_m^b).

⁹Equation (4) shows why long-run discounting only affects effort choice through $(3 - t_e)$.

Table 2. Reduced-form estimates of present bias and sophistication at the aggregate level

	Jobs	Tasks
<i>Panel A: Effort domain</i>		
ϕ_e^{future}	0.52*** (0.13)	4.17*** (1.11)
ϕ_e^{predict}	0.33** (0.15)	3.09** (1.32)
<i>Panel B: Money domain</i>		
ϕ_m^{future}	−0.05 (0.16)	−0.52 (1.34)
ϕ_m^{predict}	−0.07 (0.14)	−0.56 (1.21)
<i>Panel C: Other coefficients</i>		
ϕ_γ	0.42*** (0.01)	3.27*** (0.09)
ϕ_δ	0.12 (0.08)	1.23* (0.73)
<i>N</i>	80	80
Observations	3600	3600

Fixed effects regression estimates of Equation ((5)) from aggregate data. The dependent variable is the number of Jobs or Tasks chosen. Present bias over effort is reflected in higher values of ϕ_e^{future} , while present bias over money is reflected in lower (more negative) values of ϕ_m^{future} . The corresponding ϕ^{predict} coefficients capture perceived present bias, and the gap between ϕ^{future} and ϕ^{predict} reflects sophistication. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Note that the interpretation of coefficients differs across domains: present bias over effort is reflected in positive values of ϕ_e^{future} and ϕ_e^{predict} , whereas present bias over money is reflected in negative values of ϕ_m^{future} and ϕ_m^{predict} .

5.1.1. Aggregate analysis

Table 2 shows reduced-form measures of present bias and sophistication over the effort and money domains. As our dependent variable, we use both the number of jobs chosen, which is directly related to the payment (wage) that participants earn, and the number of real-effort tasks.

Participants choose 0.52 more jobs or 4.17 more transcription tasks (which translates to 14% more tasks¹⁰) to be completed in the future than in the present, but believe that this effect would be smaller (0.33 jobs or 3.09 tasks), which is evidence of partial sophistication. The difference between ϕ_e^{future} and ϕ_e^{predict} is not statistically significant. We do not observe a statistically significant degree of present bias in the money domain. Participants choose 0.05 fewer jobs or 0.52 fewer transcription tasks (which translates to 1% fewer tasks) when they are paid on the same day than when they are paid in the future. Neither is significantly different from zero. This result is consistent with findings by Andreoni and Sprenger (2012) and Augenblick et al. (2015), who do not find evidence of present bias over money. While the difference between ϕ_m^{future} and ϕ_m^{predict} is not statistically significant, we also do not find aggregate evidence for present bias over money in the first place.

5.1.2. Individual analysis

To derive reduced-form measures of present bias and sophistication for each participant, we re-use the regression in Equation ((5)) at the individual level for all 80 participants. Below, we only report

¹⁰The percentage is calculated by dividing the raw number by 29.39, the average number of tasks participants completed.




Fig. 4. Histograms of reduced-form estimates at the individual level. The panels on the top show estimates of *actual present bias*, while the panels on the bottom show estimates of *perceived present bias* based on prediction data. For the effort domain, actual and perceived present bias correspond to $-\phi_e^{\text{future}}$ and $-\phi_e^{\text{predict}}$, respectively, so that more negative values indicate stronger present bias. For the money domain, actual and perceived present bias correspond to ϕ_m^{future} and ϕ_m^{predict} , respectively, where more negative values indicate stronger present bias. The median and mean of actual present bias over effort are -0.2 and -0.53 . The median and mean of perceived present bias over effort are -0.14 and -0.17 . The median and mean of actual present bias over money are -0.03 and -0.06 . The median and mean of perceived present bias over money are 0.00 and -0.03 . $N = 79$ for all figures; one participant with an outlier estimate of $\phi_e^{\text{predict}} > 12.6$ is dropped

the results using the number of jobs, as the results are qualitatively identical if we used the number of tasks instead. This individual analysis is important as average values hide individual differences. As we will see, the data suggest that some individuals are present-biased, others are close to time-consistent, and others again are future-biased. Future studies can use this heterogeneity to predict theoretically related behaviours.

Fig. 4 shows the histogram of reduced-form estimates of actual and perceived present bias in both the effort and the money domains, for all but one participant, who has an outlier estimate of perceived present bias over effort of 12.6. The top-left panel shows the estimates of actual present bias in the effort domain. The $-\phi_e^{\text{future}}$ estimates of around 60% (47 out of 80) participants are below 0, which indicates a modal pattern of present bias over effort. The median is -0.2 and the mean is -0.53 .

The top-right panel shows actual present bias over money. The median and mean of ϕ_m^{future} are roughly zero: -0.03 and -0.06 , respectively. This reflects substantial heterogeneity instead of a modal pattern: 41 and 39 individuals have estimates consistent with present-bias ($\phi_m^{\text{future}} < 0$) and future-bias ($\phi_m^{\text{future}} > 0$) over money, respectively.

The bottom-left panel shows perceived present bias over effort. The median and mean of ϕ_e^{predict} are -0.14 and -0.17 , respectively. Since these measures are less negative than those of actual present bias, on aggregate, participants are partially sophisticated in our data (also see Fig. 5 for this comparison).




Fig. 5. CDFs of individual-level reduced-form estimates of actual present bias ($-\phi_e^{\text{future}}$ and ϕ_m^{future}) as represented by the lines, are compared to that of perceived present bias ($-\phi_e^{\text{predict}}$ and ϕ_m^{predict}), as represented by the dotted lines. We do this separately for effort (red) and money (blue) domains. Due to the normalization, both forms of present bias are stronger the more negative the values are. $N = 79$ for all figures; one participant with an outlier estimate of $\phi_e^{\text{predict}} > 12.6$ is dropped

The bottom-right panel shows perceived present bias over money. The median and mean of ϕ_m^{predict} are 0 and 0.03, respectively, suggesting similar heterogeneity to actual present bias over money.

The cumulative density functions presented in Fig. 5 compare individual-level reduced-form estimates of actual and perceived present bias over effort and money. Individual estimates below zero indicate the presence of actual or perceived present bias. In the effort domain, the gap between the two cumulative density functions in the top left panel shows that actual present bias exceeds perceived present bias, confirming that participants are partially sophisticated. In the money domain, the two cumulative density functions in the top right panel align closely, suggesting that actual and perceived present bias over money are similar, though we note that this need not reflect sophistication given the absence of significant present bias over money. The bottom panels compare the two domains directly: actual present bias is stronger over effort than over money (bottom left), while perceived present bias is more similar across domains (bottom right).

Fig. 6 uses four scatter plots to show the relationships between actual present bias and perceived present bias over effort and money. Again, individual estimates below zero indicate the presence of actual or perceived present bias, this time on both axes. The top left panel shows a strong relationship between actual and perceived present bias over effort (Spearman coefficient=0.72, $p < 0.0001$), suggesting substantial sophistication about present bias over effort. The further away and to the left the circles are from a 45-degree line, the less sophisticated the present bias over effort participants are. We find a strong positive relationship between actual present bias over effort and actual present bias over money (bottom-left panel, Spearman coefficient=0.75, $p < 0.0001$), suggesting that both forms of present bias are indeed related. In the bottom left panel, most of the participants ($N=33$) are in the third quadrant, i.e., are present-biased in both the effort and the money domains. We also observe more circles than cross marks, illustrating that most participants are more present-biased over effort than over money. The top right panel also shows a positive relationship between actual and perceived present bias over money (Spearman coefficient=0.46, $p < 0.0001$). We also observe




Fig. 6. Comparison of reduced-form estimates of actual and perceived present bias across domains, based on the individual-level regressions. To make the comparisons easier, observations above and below the 45-degree line are marked by circle and cross marks, respectively. The sample size is 79 for all four figures. All panels indicate the number of participants with parameter values less than 0 or greater than 0 along the respective axes

a positive relationship between perceived present bias over effort and perceived present bias over money (bottom-right panel, Spearman coefficient=0.37, $p = 0.0009$).

5.2. Structural estimation

5.2.1. Visualizing the structural identification

Before presenting the results from our structural estimation, we plot prominent features of the aggregate data that provide an intuition about how the parameters of the structural model are identified. For simplicity, we assume $\delta = 1$.¹¹ Throughout this section, we measure effort in tasks rather than jobs, as the empirical CDFs are smoother and easier to read under this scale given the non-linear mapping $e = J(J+1)/2$; all results are qualitatively identical when effort is measured in jobs instead.

¹¹Our structural estimates will confirm that this is a fair approximation for our data. Moreover, Online Appendix B.2 shows by comparing the CDFs of e_{11} and e_{22} that δ might be greater or equal to 1, without being significantly different from 1.




Fig. 7. The left figure plots immediate versus future effort and helps identify β_e . The sample size is 1305 for both future and immediate effort. The right figure plots e_{22} , $\hat{e}_{23}$, e_{33} , $\hat{e}_{22}$, and helps identify $\hat{\beta}_e$, β_m , $\hat{\beta}_m$. The sample size is 435 for all four variables

Let $e_{t\tau}$ be the effort chosen (measured either in jobs or in tasks) on day t for day τ . For a fixed decision period s , the comparison between choices for immediate effort (e_{ss}) and choices for future effort (e_{st} for $t > s$), both made at time s , identifies β_e . Under $\beta_e \leq 1$, we would observe $e_{st} > e_{ss}$ for all $t > s$, and under $\beta_e = 1$, we would observe $e_{st} = e_{ss}$. In the left panel of Fig. 7, we compare the empirical CDFs of immediate effort e_{ss} , as measured in tasks, and future effort e_{st} while pooling over wage rates and decision days ($ss = 11, 22, 33$ and $st = 12, 23$).¹² Visually, e_{ss} is first-order stochastically dominated by e_{st} , which suggests that fewer tasks are consistently allocated to the immediate day. Based on these patterns, ex-ante, we expect to estimate β_e to be significantly below 1.

Next, we explain which data patterns reveal three key parameters: $\hat{\beta}_e$, β_m , $\hat{\beta}_m$. Sophistication over effort, $\hat{\beta}_e \in [\beta_e, 1)$, is revealed if the participants predict the difference between immediate and future choices. In particular, under $\hat{\beta}_e \leq 1$, we would observe $\hat{e}_{22} < \hat{e}_{23}$, where $\hat{e}_{22}$ is the effort that on Day 1 participants predict they will choose on Day 2 to be completed on Day 2, while $\hat{e}_{23}$ is the effort that on Day 1 participants predict they will choose on Day 2 to be completed on Day 3. Under $\hat{\beta}_e = 1$ we would observe $\hat{e}_{22} = \hat{e}_{23}$. Hence, comparing $\hat{e}_{22}$, $\hat{e}_{23}$ reveals $\hat{\beta}_e$. The present bias for money, β_m , is revealed by comparing two immediate choices, e_{33} and e_{22} , where e_{33} pays on the same day and hence is not affected by β_m . Under $\beta_m < 1$, we would observe $e_{33} > e_{22}$ and under $\beta_m = 1$ we would observe $e_{33} = e_{22}$. Finally, the prediction of immediate effort $\hat{e}_{22}$ is not affected by β_e but is affected by $\hat{\beta}_m$. Thus, $\hat{\beta}_m$ is revealed by the comparison between e_{33} and $\hat{e}_{22}$. Under $\hat{\beta}_m < 1$, we would observe $e_{33} > \hat{e}_{22}$ and under $\hat{\beta}_m = 1$ we would observe $e_{33} = \hat{e}_{22}$.

¹² Assuming $\delta = 1$ allows us to also include e_{33} for the immediate versus future comparison. Not including it would still lead to identical conclusions.

Table 3. Structural estimates at the aggregate level

	[1] Same	[2] Separate	[3] No prediction	[4] Separate, $\delta = 1$
(AR2019)				
Present bias (β)	0.71 [0.52, 0.91]			
Sophistication ($\hat{\beta}$)	0.74 [0.56, 0.93]			
Present bias, effort (β_e)		0.70 [0.51, 0.90]	0.71 [0.52, 0.91]	0.79 [0.65, 0.93]
Present bias, money (β_m)		0.95 [0.64, 1.26]	0.93 [0.60, 1.27]	1.09 [0.88, 1.30]
Sophistication, effort ($\hat{\beta}_e$)		0.80 [0.53, 1.07]		0.88 [0.63, 1.13]
Sophistication, money ($\hat{\beta}_m$)		0.95 [0.66, 1.23]		1.04 [0.79, 1.29]
Discount factor δ	1.19 [1.00, 1.39]	1.10 [0.90, 1.29]	1.09 [0.90, 1.28]	
Convexity γ	2.56 [2.38, 2.66]	2.58 [2.40, 2.77]	2.58 [2.29, 2.87]	2.57 [2.39, 2.75]
Scale of effort ψ	199.36 [108.02, 290.70]	208.41 [84.17, 332.65]	211.52 [8.75, 414.29]	199.40 [86.86, 311.93]
Size of error σ	3.23 [3.17, 3.30]	3.23 [3.16, 3.30]	3.22 [3.15, 3.30]	3.23 [3.16, 3.29]
$\ln \mathcal{L}$	−6026.20	−6020.74	−4010.02	−6021.24
N	80	80	80	80
Observations	3600	3600	2400	3600

[1] *Same* follows Augenblick and Rabin (2019) to assume that present bias or sophistication over effort and money are identical. [2] *Separate* generalizes *Same* by allowing for different present bias and sophistication parameters across effort and money domains. [3] *No prediction* is similar to *Separate*, but it excludes predictions and hence cannot estimate sophistication. [4] *Separate, $\delta = 1$* imposes that the long-run discount factor is one. The numbers in square brackets indicate the lower and upper bounds of the 95% confidence interval.

To provide intuition about $\hat{\beta}_e$, β_m , $\hat{\beta}_m$, we plot the empirical CDFs of e_{22} , $\widehat{e}_{23}$, e_{33} , $\widehat{e}_{22}$ in the right panel of Fig. 7. The distributions are all close to each other and intertwined, with none first-order stochastically dominating the others. None of the three pairwise comparisons that identify $\hat{\beta}_e, \beta_m, \hat{\beta}_m < 1$, namely $\widehat{e}_{22}$ vs $\widehat{e}_{23}$, e_{33} vs e_{22} , e_{33} vs $\widehat{e}_{22}$, are significantly different. These patterns suggest that, ex-ante, we should not expect to estimate $\hat{\beta}_e, \beta_m, \hat{\beta}_m$ values that are different from 1.

5.2.2. Aggregate analysis

Table 3 presents the structural estimates at the aggregate level, using the MLE described in Section 4. For each parameter, we report the 95% confidence interval instead of the standard error, to make it easy to understand if the coefficients are significantly different from 1 or not.

The first column (‘Same’) follows the assumption from Augenblick and Rabin (2019) that present bias and sophistication are *identical* across money and effort. Thus, this specification replicates their identification assumption and serves as a counterfactual. The estimated present bias parameter is 0.71 and the 95% confidence interval lies below 1. The sophistication parameter is 0.74 and the 95% confidence interval lies below 1 as well.

The second column (‘Separate’) is our novel specification: it allows for different present bias and sophistication over effort and money. We find that β_e is 0.70 and significantly different from 1.

Interestingly, the common β estimated in the ‘Same’ specification mimics the β_e in the ‘Separate’ specification. $\hat{\beta}_e$ is 0.80, which is higher than $\beta_e = 0.70$ (although the difference is not statistically significant) and indicates partial sophistication over effort. Neither of the money-domain parameters $\hat{\beta}_m = \beta_m = 0.95$ are estimated to be significantly different from 1.

Researchers interested in estimating only the present bias parameters could use a variant of our design that does not include the prediction tasks. To approximate the results of a design with no predictions, in the third column (‘No prediction’), we run our analysis while excluding the prediction tasks (i.e., using only the data from the 6 decision screens that are not predictions). The results are very similar to the results from the ‘Separate’ specification. This is reassuring, as it indicates that including the prediction tasks does not bias the estimates of present bias.

The 95% confidence interval of δ contains 1 consistently across all three specifications. Motivated by this finding, in the fourth column (‘Separate, $\delta = 1$ ’), we re-estimate the ‘Separate’ specification assuming δ to be equal to 1 and including both choice and prediction data. While this increases the estimates of present bias over both domains, β_e is still estimated to be significantly below 1. For completeness, we note that imposing $\delta = 1$ on the restricted specification of column [1] yields a common $\hat{\beta}$ closer to 1, consistent with the pattern observed when comparing columns [2] and [4].¹³

Model comparison by Likelihood Ratio (LR) test and Akaike Information Criterion (AIC)

Under the null hypothesis that the simpler nested model in specification [1] (‘Same’) fits the data as well as the model in specification [2] (‘Separate’), the LR test statistic is 10.92, which should be distributed as chi-squared with 2 degrees of freedom. We obtain a corresponding p-value < 0.01 and thus we have statistically significant evidence to reject the null and conclude that specification [2] fits the data better than specification [1]. On the other hand, the LR test statistic in the comparison of specifications [2] and [4] is 1, which should be distributed as chi-squared with 1 degree of freedom. The corresponding p-value is 0.68 and we do not find significant evidence for model [2] fitting the aggregate data better than the simpler nested model [4]. Thus [4] fits the data almost as well as any other model without overfitting the data. Akaike Information Criterion (AIC) for model selection also suggests [4] as the superior model among [1], [2], and [4].

Stability of estimates

All estimated parameters are relatively stable across all specifications. For example, under all specifications, β_e is significantly below 1, and $\hat{\beta}_e$ and β_m are no different from 1. Similarly, the convexity parameter γ is consistently above 1, implying that the cost function is significantly convex.¹⁴ The scale of effort ψ is also similar across all specifications, implying that the baseline level of effort did not vary with the type of model chosen.

Conflated estimates from model [1]

In model [1], MLE tries to fit the misspecified model to the data by biasing δ upwards (increasing δ significantly over 1). The common β is closer to β_e than to β_m . Further, the estimated β and $\hat{\beta}$ are almost identical in magnitude under [1], implying a high degree of estimated relative sophistication. In fact, [1] is the only specification that rejects complete naïvete ($\hat{\beta} = 1$) in favour of $\hat{\beta} < 1$. As a result, participants are estimated as relatively more sophisticated under [1] than in all other specifications. For further details about how the misspecified model might generate biased estimates, see Online Appendix C.2.

¹³Online Appendix B.4 presents the summary statistics for our structural estimates. In Online Appendix B.1, we present our aggregate estimates while only dropping the participants who were in the top 5% of violations. We get very similar results.

¹⁴The convexity of the cost function guarantees that the problem in Equation ((1)) allows for a unique solution if we have an interior solution.




Fig. 8. Structural estimates of present bias and perceived present bias over effort and money. The sample size is 77 for all four graphs

5.2.3. Individual analysis

Next, we estimate the time preference parameters for each individual separately. Fig. 8 plots the distributions of the present bias and sophistication parameters, in both effort and money domains, for 77 participants.¹⁵ For present bias over effort, the distribution of β_e is skewed toward values below 1, with over 67% of participants exhibiting present bias in the effort domain, a median of 0.91, and a mean of 0.98. For four participants, β_e is less than 0.1.¹⁶ The distribution of sophistication over effort, $\hat{\beta}_e$, is similarly concentrated below 1 but less dispersed, with a median of 0.94 and a mean of 1.07, consistent with partial sophistication at the individual level. By contrast, the distributions of both β_m and $\hat{\beta}_m$ are more symmetric and centered closer to 1, with medians of 1.04 and 1.05 respectively, confirming that present bias and sophistication over money are on average weaker than in the effort domain. Still, there is substantial heterogeneity, with some participants exhibiting present bias in the money domain as well.

Fig. 9 plots the cumulative distribution functions of the individual-level structural estimates. In the effort domain (top-left panel), the distribution of present bias (β_e) lies to the left of the corresponding sophistication measure ($\hat{\beta}_e$), indicating that participants exhibit present bias but only partially anticipate it. In the money domain (top-right panel), the distributions of β_m and $\hat{\beta}_m$ largely overlap, suggesting no systematic gap between actual and perceived bias. The bottom panels compare estimates across domains: present bias is somewhat stronger for effort than for money (bottom-left),

¹⁵To make the histogram easy to read, we have removed three outliers with estimates of present bias or sophistication above 16.

¹⁶ β_e is close to 0 for participants who (i) decide to complete many more tasks in the future compared to the present and (ii) do not react strongly to changes in wage rates.




Fig. 9. Comparison of cumulative density functions of structural estimates of present bias and sophistication over effort and money, estimated at the individual level. The sample size is 77 for all four graphs

as the effort distribution is shifted toward lower values. In contrast, the sophistication distributions (bottom-right) largely overlap, providing little evidence of systematic differences across domains.

Fig. 10 presents four scatterplots of the estimated individual-level present bias and sophistication parameters in both domains. The top-left panel presents the scatter plot of β_e and $\hat{\beta}_e$. The correlation between the two parameters is high (Spearman coefficient = 0.77, $p < .0001$) and much higher than the range 0.24–0.28 reported in Augenblick and Rabin (2019) as the correlation coefficient between β and $\hat{\beta}$. The correlation suggests that participants are at least partially aware of their own present bias over effort.¹⁷ The discrepancy between our results and those of Augenblick and Rabin (2019) may be partly explained by the fact that their estimate of correlation also depends on the correlation between β_m and $\hat{\beta}_m$. The top-right panel presents the scatter plot between β_m and $\hat{\beta}_m$. The correlation between these parameters is 0.38 ($p = 0.0006$).¹⁸ The bottom-left panel shows a strong and positive relationship between β_e and β_m (Spearman coefficient = 0.77, $p < 0.0001$). Finally, the bottom-right panel presents the scatter plot of $\hat{\beta}_e$ and $\hat{\beta}_m$. The correlation is 0.30 ($p = 0.0068$).

The low correlation in the bottom-right panel is informative: knowing how sophisticated a participant is in the effort domain conveys little information about how sophisticated they are in the money domain. This contrasts with the strong correlation between β_e and β_m , indicating that underlying biases are correlated across domains even when self-awareness about those biases is not.

Further, the parametric estimates of β_e and $\hat{\beta}_e$ are strongly correlated with the reduced-form estimates of present bias and perceived present bias over effort, respectively, obtained in Section 5.1.2. Both correlations are 0.94 ($p < 0.0001$). This provides in-sample validation of the structural parameter estimates.

In the effort domain (top-left panel of Fig. 10), 39 subjects both revealed and predicted present bias ($\beta_e < 1$, $\hat{\beta}_e < 1$) while 20 both revealed and predicted future bias ($\beta_e > 1$, $\hat{\beta}_e > 1$). Thus, 59 of 77

¹⁷In Online Appendix B.3, we compare the relative sophistication of individuals by following the methodology of Augenblick and Rabin (2019). Participants understand about 26% of the present bias in the effort domain.

¹⁸The lower correlation between money domain estimates in our paper may also partly reflect lower median precision of money domain estimates in our data (see Online Appendix B.4 for details).




Fig. 10. Comparison of structural estimates of present bias and sophistication across domains. The points above and below the 45-degree line are marked with a circle and check mark respectively. The sample size is 77 for all four graphs. All panels indicate the number of participants with parameter values less than 1 or greater than 1 along the respective axes

participants correctly anticipated the direction of their bias. In the money domain, the corresponding number is 57 of 77. Table 9 in the Online Appendix further confirms this pattern by reporting the distribution of sophistication estimates conditional on whether participants were present-biased or future-biased. In both domains, participants’ predictions track the sign of their bias: those with $\beta < 1$ predict $\hat{\beta} < 1$, while those with $\beta \geq 1$ predict $\hat{\beta} > 1$. This pattern – present across roughly three quarters of all participants – points to partial sophistication in a broader sense: participants anticipate the direction of their bias, but not its full magnitude.

6. Discussion

In this section, we compare our structural estimates with three key previous studies that, like ours, directly elicit present bias and/or sophistication over effort and money (see Table 4).¹⁹ While our

¹⁹We exclude studies that discuss sophistication but do not directly measure it. For example, Zhang and Greiner (2021) use cognitive ability as a proxy for sophistication.

Table 4. Structural estimates of β and $\hat{\beta}$

	β	$\hat{\beta}$
Augenblick et al. (2015)	0.89 (β_e), 0.99 (β_m)	–
Augenblick and Rabin (2019)	0.81–0.84	1.00–1.01
Fedyk (2024)	0.82–0.86	1.03–1.05
Present study	0.70–0.79 (β_e), 0.93–1.09 (β_m)	0.80–0.88 ($\hat{\beta}_e$), 0.95–1.04 ($\hat{\beta}_m$)

key objective of estimating sophistication aligns most closely with Augenblick and Rabin (2019) and Fedyk (2024), our separate effort and money domain estimates make our study comparable with Augenblick et al. (2015).

Augenblick et al. (2015) estimate present bias separately for the effort and monetary domains in a within-subject design where participants allocated units of effort or money over two dates. For effort choices, they find significant present bias. In their main aggregate analysis, the estimated β_e is 0.89, which is statistically different from 1 ($p < 0.01$). For monetary choices, the authors fail to find evidence of present bias as the estimated β_m is 0.99, which is not statistically different from 1 ($p = 0.16$). While the authors use the concept of sophistication to link estimated present bias to observed demand for commitment, they do not measure or estimate sophistication.

Augenblick and Rabin (2019) estimate present bias and sophistication using a design where subjects make real effort decisions and predictions for different piece-rates and for different delays between the decision date and payment date. As mentioned previously, they estimate a single structural parameter for present bias, assuming it applies to both the disutility of future effort and the utility of future monetary payments. Their aggregate structural estimate for β ranges between 0.81 and 0.84 across various specifications. These estimates are significantly different from 1 ($p < 0.001$), indicating significant present bias. Across different model specifications, the estimated $\hat{\beta}$ ranges from 1.00 to 1.01. The authors cannot reject the null hypothesis that $\hat{\beta} = 1$ (p -values range from 0.11 to 0.93).

Fedyk (2024) uses the same experimental identification strategy as used in Augenblick and Rabin (2019) and extends it to also estimate participants' beliefs about others' present bias. In the main structural model, pooling across all participants, the estimated β ranges between 0.82 and 0.86. This indicates significant present bias. Fedyk documents the novel finding of asymmetric naïveté: the estimated $\hat{\beta}$ for the participants themselves ranges between 1.03 and 1.05, whereas participants perceive others' β to be 0.87.

Table 4 compares our structural estimates with those from the literature. As discussed, we deviate from Augenblick and Rabin (2019) and Fedyk (2024), as our estimation procedure does not rely on assuming the equality of present bias and sophistication across money and effort domains. This also underlines the key novelty of our design: the independent variation of the timing of effort and the timing of payments *relative to the decision date*. While all payments were received on the last day of the study, there was a variation in payment-immediacy depending on when the decision was made.²⁰ This novel double variation in task completion delay and monetary payment delay (relative to the decision point) allows us to disentangle β_e from β_m , and $\hat{\beta}_e$ from $\hat{\beta}_m$.

In our estimation exercise, while present bias over effort is significant (β_e ranges between 0.70 and 0.79), present bias over money is insignificant (β_m ranges between 0.93 and 1.09), thus justifying our identifying assumption of separate parameters.²¹ Estimated $\hat{\beta}_e$ ranges between 0.80 and 0.88 across

²⁰A decision made on Day 1 or Day 2 for a Day 3 payment is a future payment from the decision point, whereas a decision made on Day 3 for a Day 3 payment is an immediate payment from the decision point.

²¹Our finding of significant present bias over effort and insignificant present bias over money is consistent with the central takeaway from Imai et al. (2021)'s meta-analysis of 31 Convex Time Budget studies in 28 articles: the estimates for present

specifications, indicating partial sophistication over effort. Estimates for $\hat{\beta}_m$ are between 0.95 and 1.04, not significantly different from 1, though that masks significant heterogeneity.

When we repeat our estimation under the assumption of identical present bias in money and effort (Table 3, column [1]), we observe that β is very close to our estimated β_e . This might suggest that Augenblick and Rabin (2019) and Fedyk (2024)'s estimates of β , which conflate present bias over money and effort, predominantly reflect the present bias observed in effort tasks. Table 3 also suggests that enforcing $\beta_e = \beta_m$ biases δ upwards (see Online Appendix C for details).

Two factors may help explain why our participants show higher relative sophistication than those in Augenblick and Rabin (2019) and Fedyk (2024). The first relates to design differences. The shorter interval between sessions in our study – 1 day rather than 1 week – may make it easier for participants to accurately forecast their future effort choices, mechanically raising estimated sophistication. Moreover, in Fedyk (2024), *predictions about immediate effort* are elicited immediately after *choices about future effort*, which may induce them to be more similar and thus increase the gap between predictions about immediate effort versus choices about immediate effort, independently of genuine sophistication.²² In our study, participants completed the real effort tasks in between decisions and predictions, ensuring a minimal separation between decisions and predictions.

The second factor relates to the domain specificity of our estimates. Given in our study, there is little present bias to anticipate in the monetary domain, participants may appear more accurate in their predictions simply because the scope for error is smaller. More broadly, the salience of effort costs relative to monetary payments may make present bias over effort easier to anticipate.

7. Conclusion

We introduce a novel experimental design that jointly estimates present bias and sophistication across both effort and monetary domains. Our findings show strong present bias in effort but little to no present bias in money, with participants displaying partial sophistication about their own behaviour in the effort domain. This divergence across domains highlights the value of separating effort and monetary dimensions when modelling dynamic decisions.

Beyond the substantive results, our 3-day design is cost-effective, resistant to attrition, and easily replicable, making it a practical tool for applied researchers studying savings, health, productivity, or related behaviours. A natural extension would be to introduce additional variation in payment timing (e.g., allowing some same-day payments on earlier days) or to lengthen the horizon beyond 3 days. Such designs, especially in field contexts, could provide additional robustness and shed light on how time preferences shape real-world outcomes.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/eec.2026.10053>.

Replication material. The replication material for the study is available at <https://doi.org/10.5281/zenodo.19740275>.

Acknowledgements. We thank David Freeman and the audience at the VEABES for comments and feedback. An AI-based tool was used only for spelling and grammar checks. All scientific content and interpretations are entirely the responsibility of the authors.

Funding statement. A.C. thanks UC Davis Small Grants and Global Affairs Seed Grant for financial support.

bias under monetary rewards are close to one, indicating the absence of present bias, while the same from studies with non-monetary rewards report a lower average of 0.88, implying significant present bias. One of our anonymous referees suggested an interesting hypothesis that people pay more attention to the date effort needs to be exerted than to the date payments are made. This could help explain why present bias appears primarily in the effort domain. While our design cannot directly test this mechanism, it highlights a promising direction for future research.

²²We thank an anonymous referee for suggesting this. We also note that confidence intervals across studies are wide, so some apparent differences may reflect sampling variation rather than true design effects.

References

- Ahn, D. S., Iijima, R., Le Yaouanq, Y., & Sarver, T. (2019). Behavioural characterizations of naivete for time-inconsistent preferences. *The Review of Economic Studies*, 86(6), 2319–2355.
- Andreoni, J., & Sprenger, C. (2012). Estimating time preferences from convex budgets. *American Economic Review*, 102(7), 3333–3356.
- Ariely, D., & Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13(3), 219–224.
- Ashraf, N., Karlan, D., & Yin, W. (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *The Quarterly Journal of Economics*, 121(2), 635–672.
- Augenblick, N., Niederle, M., & Sprenger, C. (2015). Working over Time: Dynamic inconsistency in real effort tasks. *The Quarterly Journal of Economics*, 130(3), 1067–1115.
- Augenblick, N., & Rabin, M. (2019). An experiment on time preference and misprediction in unpleasant tasks. *Review of Economic Studies*, 86(3), 941–975.
- Carrera, M., Royer, H., Stehr, M., Sydnor, J., & Taubinsky, D. (2022). Who chooses commitment? Evidence and welfare implications. *The Review of Economic Studies*, 89(3), 1205–1244.
- Chakraborty, A. (2021). Present bias. *Econometrica*, 89(4), 1921–1961.
- Chakraborty, A., Calford, E. M., Fenig, G., & Halevy, Y. (2017). External and internal consistency of choices made in convex time budgets. *Experimental Economics*, 20(3), 687–706.
- Chakraborty, A., & Fenig, G. (2025). From me to we: Beating procrastination in teams *Journal of Economic Behavior & Organization*, 231, 106900. <https://doi.org/10.1016/j.jebo.2025.106900>.
- Cheung, S. L., Tymula, A., & Wang, X. (2022). Present bias for monetary and dietary rewards. *Experimental Economics*, 25(4), 1202–1233.
- Cobb-Clark, D. A., Dahmann, S. C., Kamhöfer, D. A., & Schildberg-Hörisch, H. (2024). Sophistication about self-control. *Journal of Public Economics*, 238, 105196.
- Ericson, K. M., & Laibson, D. (2019). Intertemporal choice. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics – Foundations and applications 2* (vol. 2, pp.1–67). North-Holland.
- Fedyk, A. (2024). Asymmetric naivete: Beliefs about self-control. *Management Science*, 70(5), 2633–2657.
- Gerhardt, H., & Suchy, R. (2024). Estimating preference parameters from strictly concave budget restrictions. Technical report, ECONtribute Discussion Paper.
- Gul, F., & Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69(6), 1403–1435.
- Imai, T., Rutter, T. A., & Camerer, C. F. (2021). Meta-analysis of present-bias estimation using convex time budgets. *The Economic Journal*, 131(636), 1788–1814.
- Kaur, S., Kremer, M., & Mullainathan, S. (2015). Self-control at work. *Journal of Political Economy*, 123(6), 1227–1277.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2), 443–478.
- Le Yaouanq, Y., & Schwardmann, P. (2022). Learning about one's self. *Journal of the European Economic Association*, 20(5), 1791–1828.
- O'Donoghue, T., & Rabin, M. (1999). Doing It Now or Later. *American Economic Review*, 89(1), 103–124.
- Phelps, E. S., & Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies*, 35(2), 185–199.
- Sadoff, S., & Samek, A. (2019). Can interventions affect commitment demand? A field experiment on food choice. *Journal of Economic Behavior & Organization*, 158, 90–109.
- Schilbach, F. (2019). Alcohol and self-control: A field experiment in india. *American Economic Review*, 109(4), 1290–1322.
- Sial, A. Y., Sydnor, J. R., & Taubinsky, D. (2023). Biased memory and perceptions of self-control. Technical report, National Bureau of Economic Research.
- Sprenger, C. (2015). Judging experimental evidence on dynamic inconsistency. *American Economic Review*, 105(5), 280–285.
- Strack, P., & Taubinsky, D. (2026). Dynamic preference “reversals” and time inconsistency. *American Economic Journal: Microeconomics*, 18(2), 426–468.
- Toussaert, S. (2018). Eliciting temptation and self-control through menu choices: A lab experiment. *Econometrica*, 86(3), 859–889.
- Westphal, R. (2024). People do not demand commitment devices because they might not work. *Journal of Economic Behavior & Organization*, 228(C), 106756.
- Zhang, Q., & Greiner, B. (2021). Time inconsistency, sophistication, and commitment: An experimental study. *Economics Letters*, 206(C), 109982.