Eliciting time preferences when income and consumption vary: Theory, validation & application to job search

Michèle Belot, Philipp Kircher, and Paul Muller*

May 2, 2024

Abstract

We propose a simple method for eliciting individual time preferences without estimating utility functions even in settings where background consumption changes over time. It relies on eliciting preferences for receiving high stakes lottery tickets at different points in time. In a standard intertemporal choice model high rewards decouple lottery choices from variation in background consumption. We investigate robustness to other assumptions theoretically, and validate our elicitation method experimentally. We illustrate an application of our method with unemployed job seekers which naturally have income/consumption variation.

Keywords: time preferences, experimental elicitation, job search, quasi-hyperbolic discounting
JEL-codes: D90, J64

* *Affiliations*: Michèle Belot (mb2693@cornell.edu), Cornell University and University of Edinburgh; Philipp Kircher (pk532@cornell.edu), Cornell University, Catholic University of Louvain and University of Edinburgh; Paul Muller (p.muller@vu.nl), Vrije Universiteit Amsterdam and Tinbergen Institute. This study was pre-registered on the AEA RCT Registry, AEARCTR-0011729. We are grateful to Lucia Casal for her excellent research assistance, and thank many seminar audiences for their valuable suggestions. We are particularly grateful to Glenn Harrison, Ted O'Donoghue, Andrew Clausen, Karlijn Morsink and Coen Teulings. Kircher gratefully acknowledges support through European Research Council Grant No 284119. All data used in this article are accessible from Belot et al. (2024).

1 Introduction

Recent research has increasingly acknowledged the importance of behavioral elements, such as time preferences and beliefs about job finding, in the context of job search (e.g. DellaVigna and Paserman (2005), DellaVigna et al. (2017), Dohmen (2014), Caliendo et al. (2015), Spinnewijn (2015), Mueller et al. (2021)). Understanding their role is important for informing policy design. In this paper we highlight the challenge of measuring time preferences in a context such as job search, where income and leisure are likely to vary substantially over time.

Experimental elicitation methods based on monetary rewards typically evaluate the marginal utility of money at different points in time. In a context where expectations about consumption levels vary across periods, such elicitation methods will confuse the discount factor with these expectations. Concretely, unemployed individuals tend to have limited income currently but more income once they find a job. There is evidence that their consumption varies substantially with income, particularly for job seekers with low levels of education and low levels of assets (Dynarski and Gruber (1997), Ganong and Noel (2019)).[1] In such a setting, a job seeker who expects to find a job quickly would prefer money today because she has less need for it in the future, while an otherwise identical job seeker who expects that she will have to search long to find a job is more willing to take money in the future as she also expects to be poor then.[2] Without adjusting for the expected change in income, the former would appear more impatient than the latter. And it would look as if the more impatient person finds a job quicker with less effort, while in fact both have the same impatience but different (possibly true) expectations about job finding. Job search seems a particularly stark setting where income and consequently consumption change rapidly over time, but variation across seasons in agricultural economies or settings with uninsured risk to income or expenditure (e.g., medical costs) raise similar concerns.

The researcher has several options to deal with changing background consumption at the individual level. She could elicit expectations about future income and consumption, and structurally back out a discount factor. The drawback of this approach is that she would have to elicit a large object: the probability of finding a job at various time horizons, the expected distribution of starting wages, and possibly wage progression and job loss probabilities.[3] While constant average income and associated background consumption have been important controls in influential studies (e.g., Andersen et al. (2008), Andreoni and Sprenger (2012)), we are not

[1] Dohmen (2014) emphasizes that theoretically the implications of deviations from exponential discounting for job search have been derived, but convincing empirical evidence is still scarce and requires development of better elicitation methods.

[2] The important point about this thought experiment is that both individuals have the same true discount factor and the same utility function. The former may well have higher unobserved skills that are more desirable in the labor market and rationalizes her expectations relative to the second individual.

[3] Job loss probabilities are not irrelevant, as many jobs last only for very short spells.

aware of studies that directly deal with varying background consumption at the individual level.

The researcher could draw on elicitation methods that rely on real-effort, i.e., that are not in the monetary domain (e.g., Augenblick et al. (2015)). In the case of job search unfortunately also the endowment of free time and arguably of free mental resources changes with employment status. A person who expects to find a job quickly might prefer to do a real effort task now while she still has free time (and, thus, appearing patient) while a person who expects to be unemployed for longer might prefer to wait, even if their true discount factors are identical. In other words, since both background consumption and leisure are likely to vary sharply over time, none of the methods available - whether based on money or real effort - can robustly elicit time preferences.

In this paper, we build on the idea of methods that aim to elicit individual's time preferences "directly", i.e., without estimating utility functions. The goal here is a first pass at this problem via a simple change to standard methods that deals with income and associated background consumption changes without the need to explicitly back out the background consumption for each individual. We propose a simple monetary elicitation method based on high stakes lottery tickets. It is based on a straightforward intuition: if one asks individuals whether they prefer to be "rich" early but at lower probability or later but with higher probability, the exact nature of their current income/consumption stream should not matter as long as being "rich" is sufficiently far removed from their usual income/consumption. We first show that this logic indeed holds theoretically in the simplest model. We then extend it along several dimensions: savings, probability weighting, and separation of intra- and inter-temporal curvature in utility. Our approach combines two innovations: 1) making rewards probabilistic, and 2) making rewards in some states (after winning the lottery) large. The second allows for the separation from background consumption, while the first offers a number of practical advantages including enabling this approach without a massive research budget, without a need to back out individual utility parameters, and without a need to worry that individuals borrow against future rewards if the probabilities are sufficiently small.

In terms of implementation, we rely on an elicitation method that resembles the standard Multiple Price List approach, only that individuals do not decide about money but about the number of high stakes lottery tickets. Each represents the probability of winning a large prize. Given a probability of winning early, for each individual we can elicit a probability of winning late that makes the person indifferent. In our baseline model the ratio of the early and the equivalent late winning probability approximates the true discount factor, both in the short and in the long run, as long as the reward conditional on winning is large. We also present extensions allowing for rewards to be consumed over multiple periods mechanically or through optimal savings, as well as allowing for probability-weighting and for separation of intra-temporal risk aversion and inter-temporal elasticity of substitution of consumption. While the level of the discount factor cannot always be uncovered, we show that our method can back out robustly the ranking of individuals: whether one individual is more patient over a particular horizon than another. This ranking is relevant if the goal is to correlate patience in the short or long run with other outcomes, such as job search effort. The fundamental insight is that in all our theoretical specifications, measurement is not influenced by the level or the changes in other income and consumption. This constitutes the main innovation: the identification of the level/rank of patience across individuals that have varying and heterogeneous background consumption.

Asking individuals about the probability and timing of winning a large prize might entail its own cognitive challenges for individuals. In our implementation we attempt to make the choice cognitively easy by asking individuals about the number and timing of scratch-off lottery tickets. What matters according to the theory is the ratio of winning probabilities. So as long as individuals perceive, e.g., 2 lottery tickets roughly twice as likely to win as 1 lottery ticket, the results apply.[4] Individuals do not need to know the actual level of the probability.

Despite this reassurance, it might be useful to evaluate our method experimentally. After all, our approach relies on the idea that individuals consider the timing and probability of becoming "rich" in a way that abstracts from fluctuations in their income and consumption streams. In standard economic theory, this remains tightly linked to the discount factors and decisions that individuals take at normal stakes at least in the absence of fluctuations, but it remains an empirical question whether this is indeed the case. We therefore investigate whether our method is robust to income shocks by recruiting 305 participants on the Prolific platform. We elicited discount factors for participant using both the common Convex Time Budget (CTB) method of Andreoni and Sprenger (2012) and our new Multiple Lottery List (MLL) method.

Since actual variations in income and background consumption are unlikely to be exogenous, we design an intervention that aims at generating a variation in *perceived expenditures*. Specifically, we ask participants to report an estimate of their expected expenditures for the next 3-4 weeks (as one number) and introduce a treatment where a subset of individuals is asked to provide details of these expenditures: rent & mortgage payments, property tax payments, auto insurance premium, out-of-pocket medical costs, life insurance premiums, electricity and natural gas bill, water bill, credit card payments, extraordinary repairs and the likes. It is long established in psychology and economics that salience matters (e.g., Taylor and Thompson (1982); Bordalo et al. (2013)) and explicitly enumerating expenditures has been shown to affect individuals future financial decisions (Augenblick et al. (2022)). As we anticipated in the

[4]We acknowledge that obtaining two lottery tickets could involve the chance of winning the prize twice if the draws are independent, but we show that this does not substantially alter the insights as long as the winning probability is sufficiently small.

pre-registration, relative to a blanket question about expenditure, we see that these individual questions (which we refer to as "prompting") indeed raise the total costs that these individuals perceive, especially in the short-run. For things to add up, they either have to consume less, or adjust their savings, akin to a short-term income loss.

Our main hypothesis was that this does not affect measured present-bias in our method. If these income changes are strong enough, they should raise measured present bias in traditional methods. We introduced this variation both within and between individuals: the within-individual method has more statistical power, but we included the between-individual variation due to potential 'anchoring' when preferences are elicited twice for the same individuals. So some individuals answered MLL (or CTB) questions both before and after prompting, but we can also compare the answers before prompting to answers from other individuals who are only asked after prompting.

Reassuringly, if these questions are asked prior to the intervention, the discount factors elicited with MLL are significantly correlated with those under CTB. This applies both to the long-term discount factor and the measure of present-bias. So these measures seem to align well, despite the probabilistic nature and different payoff domain in MLL.

Considering the more pressing question of whether our discount elicitation method indeed protects against income shocks, we find that the 'prompting' treatment does indeed not affect elicited measures of time preferences with our proposed method (MLL). Under CTB, we do not find significant treatment effects between individuals, but we find that the treatment significantly increases present-bias (the average estimated beta decreases from 1.03 to 0.98, p-value 0.049) in the within-treatment. We also pre-registered an exploratory analysis to understand the stability of preferences in the within design: we regress the pre-prompting estimate of an individual on the post-promption estimate, and find that these are very highly significantly correlated under MLL while they are not significantly related under CTB.[5] This property of MLL might allow this method to address discrepancies between time consistency and time stationarity that have been pointed out in the literature, as we elaborate when discussing the related literature.

Finally, we return to our initial motivation and illustrate how our method can be deployed in the context of our leading example: job search. Since unemployed individuals naturally experience varying income streams associated with re-employment, we initially developed our method as a way to elicit discount factors with this population (see Belot et al. (2019b) for experimental details). It preceded the full development of all parts of the theory presented here, and the validation exercise.

[5]The regression coefficients are significantly different, with the MLL coefficient being much higher than the CTB coefficient, which was the main hypothesis listed for exploratory analysis. The latter turns out to be close to zero and insignificantly different from zero.

We find a bi-modal distribution of long-run discount factors. Most unemployed are not present-biased, and only 20% can be characterized as such. Those with less present-bias receive more job interviews, have positive (but not significant) increases in job search effort, with no impact on reservation wages, controlling for standard observables. A higher long run discount factor is associated with less job interviews, less job search, and no impact on reservation wages. The first is broadly consistent with predictions of DellaVigna and Paserman (2005), while the second suggests that additional forces are at play and we briefly discuss potential candidates. We discuss related literature in the next section, but are not aware of studies that aim to address this with a method designed to deal with changing background consumption.

Fluctuations in income/consumption are not limited to the unemployed. They might equally apply to farmers in developing countries, or to small entrepreneurs. Even average US workers face substantial earnings risk over time (e.g., Guvenen et al. (2021)). Our method offers a number of advantages in eliciting discount factors in such settings: It is relatively easy and fast for participants in an experiment, especially if they can just choose how many lottery tickets to have early vs late. This is why we chose it in the job search experiment. It does not need much information about the utility function. Nor does it need to make assumptions about the exact timing or nature of income fluctuations. Its disadvantages include the use of questions that might be unusual for participants, and hurdles for the experimenter if appropriate lottery tickets are not commercially available. A possible alternative is to use existing discount factor elicitation's methods coupled with questions about current and future consumption expectations. We are not aware of a "direct" method that could convert these into discount factors without the need to recover each individual's utility function. Methods that jointly determine an individual's utility function and discount factor such as CTB would have to be structurally estimated. The benefit would be to work in a framework that is well-understood in the absence of income fluctuations, but researchers not only have to take a stance on the class of utility functions but also on the number of moments of future earnings/consumption that need to be elicited from participants. The literature on earnings dynamics following Guvenen et al. (2021) argues that the mean does not suffice and higher moments up to kurtosis are highly relevant for earnings risk and earnings dynamics. Also different persistence of positive and negative earning shocks are important, indicating that belief elicitation and structural estimation may be a non-trivial exercise. Such a structural approach would be particularly challenging for research questions that aim to study the link between discount factors and expectations about future consumption, since the latter is also structurally needed to estimate the discount factors.[6] We are not aware of existing attempts

[6]For example, if one wants to correlate discount factors with job search effort and expected job finding rates (and associated income and consumption increases), it is conceptually challenging to use measures of expected job finding rates (or associated consumption increases) in the estimation of the discount factor: any errors in the measure of expectations would induce systematic errors in the measure of the discount factor, correlating them

in this direction, but believe that such developments would be complementary to the elicitation method that we propose here.

The rest of the paper is structured as follows. Secion 2 discusses the most related work. Section 3 presents the theory in the simplest model by building on classical methods and expanding to ours, and Section 4 explores robustness to consuming rewards over multiple periods either mechanically or through optimal savings, to probability-weighting and to separation of intra-temporal and inter-temporal curvature in utility. Section 5 presents the validation experiment, and Section 6 presents our discount factor measures and correlations with job search activities in a sample of job seekers. Section 7 concludes.

2 Related literature

The challenge of eliciting time preferences with varying background consumption is relatively well-known in the literature, but solutions are very scarce. The elicitation methods such as Multiple Price Lists (MPL hereafter, Coller and Williams (1999), Harrison et al. (2002)), risk and time preference integration (Andersen et al. (2008)), or the "direct" methods that aim to estimate time preferences without utility estimation (Laury et al. (2012); Attema et al. (2016)) cannot account for individual changes in background consumption, or would require eliciting precisely how income and expenditures is expected to change over the time at the individual level. Even those methods where utility terms cancel and time preferences are directly recovered require constant background consumption, otherwise the cancellation argument does not apply.

A number of empirical studies point out that variations in income may contaminate the estimates of time preferences. Ambrus et al. (2014) conducted an experiment in Iceland where they delayed rewards on actual income realizations (observed using administrative tax records). They find that background income is associated with preferences over unconditional delayed rewards, and that while most people exhibited present bias when comparing unconditional delayed rewards, subjects with stable income did not. Carvalho et al. (2016) compare individuals in the US before and after payday and find the before-payday group behaved as if they were more present-biased when making inter-temporal choices about monetary rewards. Very recently, Dean and Sautmann (2020), using data from a panel study in Mali, point out that the measures of time preferences cannot be interpreted as such in the presence of income shocks but can be used to inform about credit constraints individuals face.

Further evidence relates to how credit and liquidity constraints may contaminate the elicitation of time preferences. Pender (1996) measured discount rates in rural India (using MPL), and argued that exorbitantly high real interest rates charged by moneylenders generate binding

in a spurious way. We thank an anonymous referee for alerting us to this potential issue in structural models.

credit constraints and explain individuals' excess impatience. Holden et al. (1998) find that liquidity-constrained households in rural areas in Indonesia, Zambia and Ethiopia exhibit much higher discount rates than households not facing such constraints. Meier and Sprenger (2010) correlate experimentally elicited measures of time preferences (using MPL) with data on credit borrowing and find that more present biased individuals are more likely to have debt.

Andreoni and Sprenger (2012) discuss the issue of changing background consumption in their proposal to use convex budget sets for elicitation. They propose a empirical model to estimate quasi-hyperbolic discount rates. Using their experimental data, they back out two additional parameters for current and future background consumption, in a version of the model where discount rates are identical across individuals. When estimating individual discount rates, they place restrictions on the values for current and future background consumption. Here the concern is complementary, by focusing on variation at the *individual* level, where discount factors and background consumption might differ for each individual. This is particularly relevant in the literature in labor and development economics that intends to uncover the true discount factor, and then to correlate this with other variables such as individuals' job search behavior or borrowing (such as Meyer (2018) or Meier and Sprenger (2010)). So it is of particular interest to uncover which individual is more patient than another, i.e., to rank individuals by their patience, rather than to uncover the absolute level of a population-wide discount factor.

Our method is closest to Laury et al. (2012) who propose to elicit time preferences using binary lottery choices to avoid having to estimate risk aversion parameters separately and thereby conveniently "avoid the curves". Their method still only applies to settings with stable background consumption. Gerber and Rohde (2015) extends that setup to settings with changing background consumption, and prove theoretically how one can test for constant, increasing or decreasing discounting over time without estimating the utility function. Only if the utility function of each individual and their early background consumption can be structurally estimated, then theoretically a second elicitation can be used to back out the additional change in consumption, and then the discount factor, though only if there is no uncertainty in background consumption.[7] These papers consider lotteries with moderate rewards conditional on winning, while we consider very large rewards. As we show, the size of rewards is crucial to back out the discount factors directly without knowledge of the utility function, even with time-varying and potentially uncertain background consumption, as it allows us to cancel utility terms in the limit.

Note that despite considering large potential rewards, our environment still constitutes a

[7]It is not clear how this would survive with uncertainty about background consumption, as this would require knowledge of all differences in background consumption between the early and late period, which tends to require either elicitation of all these, or many indifference conditions according to their method.

"small stakes" environment because the *expected* payments in our setting remain small. Just conditional on winning our rewards are large. Work that has considered large stakes per se has shown a magnitude effect whereby higher rewards lead to lower measured discount factors (e.g., Andersen et al. (2013)). This has not yet been explored as a solution to time-varying background consumption.

Nearly all methods for eliciting discount factors are derived from models with additively separable utility. We follow this tradition. A notable exception is Andersen et al. (2018) who separate intratemporal risk aversion and intertemporal elasticity of substitution to capture correlation aversion with respect to stochastic rewards. In the theoretical extensions we study how our procedure performs in their environment without time-separable preference.

Real effort elicitation methods (Augenblick et al. (2015)), which have recently increased in popularity, remain vulnerable to changes in time endowment (akin to changes in background consumption) as discussed above, and we did not see a similar way to deal with this through large rewards akin to our implementation in the monetary domain.

Much work on present-bias has focussed on time stationarity: when asked today, does a decision-maker value consumption today over consumption in a month differently from consumption in a month vs consumption in two months. Recent work has focused on the separate dimension of time consistency: whether a decision-maker values immediate rewards today over rewards in a month similarly or differently when he/she is asked the same question again in a month. These dimensions are distinct, and Halevy (2015) shows that only three quarter of those who are time consistent are also time stationary, and about half of those who are not time consistent still exhibit stationary preferences. He attributes this to background consumption risk: immediate background consumption is known but the future is uncertainty, and consumption in the future might be more or less constraint than is anticipated today. Our MLL method is intended to limit the role of background consumption for time preference elicitation, and therefore might offer a way to resolve these empirical inconsistencies. The exploratory finding in our experiment that measured present-bias before prompting individuals to worry about their expenditures are highly correlated with those after prompting under MLL (but not under CTB) offers some first indication in this direction, though we do not actually repeat experimental questions over extended periods of time.

Finally, there are very few other studies that elicit direct measures of time preferences from the unemployed and relate them to job search, maybe because of the issues that our method is intended to resolve. A recent study by Le Bihan (2023) correlates measures of time preferences with measures of job search in a sample of French unemployment job seekers. He uses different 'conventional' methods to elicit preferences (Double Multiple Price List and CTB, in the money and effort domains). The correlations appear to be very sensitive to the method used. Hall and

Mueller (2018) included a simple choice between \$40 reward to fill their survey early or \$20 to fill their survey later, and find no difference in reservation wages between those who chose early vs those who choose late. Meyer (2018) applied the method of convex budget sets of Andreoni and Sprenger (2012) to study job search of employed individuals in a developing country. For employed individuals income changes might be less pronounced than for the unemployed. He reports lower job search and job finding for those with higher measures of present bias.

3 Time preference measurement: A Simple Model

Here we consider a very stylized setting for expositional simplicity. Consider an individual i who cannot save or borrow, and is an expected utility maximizer. Both of these simplifications will be relaxed later (see Section 4). The individual faces exogenous and potentially stochastic net income stream $y_{i,t}$ in periods $t = 0, 1, ..., T$, where T could be infinite. Net income comprises income in this period $(Y_{i,t})$ minus unavoidable costs in this period $(C_{i,t})$ such as tax payments, school fees, unexpected medical bills or unavoidable repairs. Assume possible net income realizations are in a bounded subset $[\underline{y}, \bar{y}]$. In period t the individual obtains consumption utility $u_i(c_{i,t})$ from consuming $c_{i,t} = y_{i,t} + r_{i,t}$, where $r_{i,t}$ denote any additional rewards that arise as part of the experimental setup and $u_i(\cdot)$ is a twice-differentiable, strictly increasing and strictly concave Bernoulli utility function. Similar to many other papers, this simple setting assumes that individuals are "hand-to-mouth", i.e., they consume what they earn. Therefore, the term $y_{i,t}$ is also often called "background consumption". We assume throughout that the utility function is unbounded and has vanishing first derivatives (i.e., $\lim_{c\to\infty} u_i(c) = \infty$ and $\lim_{c\to\infty} u_i'(c) = 0$).

The individual maximizes the expected discounted sum of the consumption utilities, where $\gamma_{i,\tau,t}(t_0)$ denotes the discount factor for an individual at time t_0 who chooses between an earlier period $\tau \geq t_0$ and a later period $t > \tau$. Unless otherwise noted we normalize $t_0 = 0$ and omit it from the notation. This particularly simple setup is similar to those used to inform other methods of discount factor elicitation and serves as illustration. We later allow for richer environments in the robustness section.

The discount factor $\gamma_{i,\tau,t}$ is the object we want to elicit. Following the literature, we assume multiplicative separability: $\gamma_{i,\tau,t} = \prod_{s=\tau}^{t-1} \gamma_{i,s,s+1}$. Particularly well-known cases are those of exponential discounting, where $\gamma_{i,s,s+1} = \delta_i$ for some per-period discount factor $\delta_i \in (0, 1)$, or *quasi-hyperbolic* discounting that has an identical structure as exponential discounting for future choices $(s > 0)$ but uses $\gamma_{i,0,1} = \beta_i\delta_i$ with $\beta_i < 1$ when the early outcome involves immediate gratification $(s = 0)$.

Standard Approach: an illustration using Multiple Price List (MPL). To investigate the discount factor $\gamma_{i,\tau,t}$, assume that the individual is offered a choice between r_τ units of additional money (and consumption) in an earlier period τ or r_t units of additional money at the later time t. Multiple Price Lists (e.g., Coller and Williams (1999) and Harrison et al. (2002)) present the individuals with varying amounts of - say - the late reward, where only one of the choices is eventually selected for payout. Fixing r_τ and evoking expected utility theory, the point $r^*_{i,t}$ below which the individual prefers early rewards and above which the individuals prefers late rewards marks the point of indifference. Since no other periods are affected, indifference arises when

$$Eu_i(y_{i,\tau} + r_\tau) + \gamma_{i,\tau,t} Eu_i(y_t) = Eu_i(y_{i,\tau}) + \gamma_{i,\tau,t} Eu_i(y_{i,t} + r^*_{i,t}), \tag{1}$$

where E represents her income expectations, which can depend on individual i and period t without making this explicit. The left hand side of (1) is the expected utility of consuming the income and reward early plus the discounted future consumption of income. At the point of indifference it should coincide with the right hand side where a larger reward is paid out after discounting. This can be rearranged to yield

$$\gamma_{i,\tau,t} = \frac{Eu_i(y_{i,\tau} + r_\tau) - Eu_i(y_{i,\tau})}{Eu_i(y_{i,t} + r^*_{i,t}) - Eu_i(y_{i,t})} \approx \frac{r_\tau}{r^*_{i,t}} \frac{Eu'_i(y_{i,\tau})}{Eu'_i(y_{i,t})}, \tag{2}$$

where the approximation is valid for small rewards when a Taylor approximation $u_i(y + r) \approx u_i(y) + u'_i(y)r$ can be applied. It led to the use of the reward ratio $r_\tau / r^*_{i,t}$ as a measure of time preference, which is valid either if utility is essentially linear or if expectations over consumption streams are roughly constant over time. Influential work has pointed out that an approximation as in the second part of (2) might not be valid and there are ways of adjusting for this (Anderson et al (2006), Andreoni and Sprenger (2012)). The main point here is that even if it were valid, there are concerns about discount factor measurement when incomes vary substantially over periods.

The problem in our setting: If utility is strictly increasing and strictly concave, it is easy to highlight the bias alluded to in the introduction. Consider a job seeker with low income currently (period τ) that expects an increase in earnings in two months (period t) due to finding a job. With standard parameter choices we can calibrate the individual's choices in a multiple price list experiment and the resulting estimate for their discount factor.[8] If their true discount factor is 0.993 for the two months (equivalent to 0.96 annually), the expected increase in income increases the experimental amount required for the individual to pick the later option. As a result, the estimate for the discount factor would be 0.46, far lower than its true value. As we show in

[8] We use log-utility, current income equal to 668 dollar per month and future income equal to 1440 dollar per month. These values are taken from our job seeker study in Edinburgh (see section 6). In the Online Appendix B we show that the results hold regardless of any of the parameter choices.

Online Appendix B, this bias occurs also with relatively smaller variations in future income.

The same concern that we outlined here with money applies with real effort if individuals without a job have a lower marginal utility of time than those who have a job. Online Appendix A.1 illustrates this. A similar problem occurs in the direct method (Attema et al. (2016)) that is designed to identify the discount factor without the need to identify the underlying utility functions if income/background-consumption is steady, but is affected by changes in the marginal utility of money when income/background-consumption can change (see Online Appendix A.2).

Alternative: Multiple lottery lists (MLL). We aim for an elicitation method that does not vary with the expected income stream of the individual and her associated background consumption, and does not require knowledge of each individual's utility function. In essence, we aim to find a method in which a simple statistic, similar to the reward ratio in MPL, elicits the discount factor. To achieve this, we exploit two basic ideas. One is to exploit the advantage of a monetary elicitation method in which - in principle - one could make the experimental rewards very large. It can be shown that even in the MPL method this decouples the reward ratio from varying background consumption. It does have two drawbacks, though: With large rewards the discount factor is no longer represented by the ratio of reward, but by the utilities that each individual obtains from these rewards.[9] More concerning, in most settings the experimenter does not have the funds to cover such high rewards. Therefore, the second innovation in this paper is to use lottery tickets that are commercially available at affordable prices and pay out the reward only with some probability.

This leads us to the following alternative elicitation method where the individual returns are probabilistic, i.e., where participants receive a large amount R of additional income with some small probability, akin to a lottery win. They can choose to get some number n_t of lottery tickets in the early period, or a larger number n_τ of lottery tickets in the later period. If the winning probability per ticket ϵ is low, than the chance of a win is roughly $\varepsilon_{tau} = n_\tau \epsilon$ for the early choice and $\varepsilon_t = n_\tau \epsilon$ for the later choice. Important for the implementation is not the level of these probabilities, but that individual's understand that the latter is roughly n_t/n_τ times more likely to win. In the case of $n_t = 2$ and $n_\tau = 1$ that means that individuals understand that the two later lottery tickets are roughly twice as likely to lead to a win compared to the one early lottery ticket. In the implementation we choose lottery tickets that can be "scratched"

[9] In Online Appendix A.3 we show that under the MPL method the reward ratio $r_\tau/r^*_{i,t}$ becomes unaffected by background consumption if reward r_τ grows large. Yet, this reward ratio no longer represents the discount factor, since the discount factor now equals the ratio of the utilities $u_i(r_\tau)/u_i(r^*_{i,t})$. This arises because the second approximation in (2) - which eliminated the utility function from the equation under the assumption of small rewards and constant background consumption - is no longer valid. So moving from small to large stakes eliminates concern for background consumption, but requires to depart from using the reward-ratio and to know the properties of the individual's utility function.

off and immediately redeemed. The idea is here is that conditional on winning the return is so large that the exact levels of other income do not matter any longer.

Formally, we start this derivation by considering an individual that can either obtain a large additional income R with probability ε_τ in the earlier period τ, or with probability ε_t in the later period t. We assume that individuals can only win the reward R once, but treat this and the exact implementation with lottery tickets more carefully in the section, together with other important extensions. Fixing the early probability ε_τ, we vary the later winning probability akin to the approach under Multiple Price Lists. The point where the individual switches from preferring to win early to preferring to win late marks the winning probability $\varepsilon^*_{i,t}$ of indifference where

$$\begin{aligned} & [(1-\varepsilon_\tau)Eu_i(y_{i,\tau}) + \varepsilon_\tau Eu_i(y_{i,\tau}+R)] + \gamma_{i,\tau,t}Eu_i(y_{i,t}) \\ = \quad & Eu_i(y_{i,\tau}) + \gamma_{i,\tau,t}[(1-\varepsilon^*_{i,t})Eu_i(y_{i,t}) + \varepsilon^*_{i,t}Eu_i(y_{i,t}+R)], \end{aligned} \tag{3}$$

where, for simplicity, we denote by E the expectations with respect to normal income and explicitly account for the uncertainty about the lottery win R. Rearranging leads to

$$\begin{aligned} \gamma_{i,\tau,t} \quad &= \quad \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \frac{Eu_i(y_{i,\tau}+R) - Eu_i(y_{i,\tau})}{Eu_i(y_{i,t}+R) - Eu_i(y_{i,t})} \\ &\approx \quad \underbrace{\varepsilon_\tau/\varepsilon^*_{i,t}}_{\textit{probability ratio}} \qquad \text{for large } R \text{ or } Eu_i(y_{i,t}) \approx Eu_i(y_{i,\tau}). \end{aligned} \tag{4}$$

The approximation is obvious when income expectations are similar across periods. More importantly, when the reward R conditional on winning becomes sufficiently large, similar income expectations are not longer necessary, which is the entire point of this exposition. This is obvious as the second ratio on the right hand side of converges to $(u_i(R) - Eu_i(y_{i,\tau}))/(u_i(R) - Eu_i(y_{i,t}))$ because of vanishing marginal utilities, and this is bounded between $(u_i(R) - u_i(\bar{y}))/(u_i(R) - u_i(\underline{y}))$ and $(u_i(R) - u_i(\underline{y}))/(u_i(R) - u_i(\bar{y}))$, both of which converge to 1.

Therefore the *probability ratio* identifies the discount factor for large R (as in the case of actual lottery tickets), and the reliance on the conditions of the regular income process disappears when the utility function is unbounded, allowing a measurement of the discount factor through the adjustment in probabilities. Note that this persists even if we keep the expected value of the early lottery fixed, i.e., $\varepsilon_\tau R = K$ for some constant K. In that interpretation the probability ratio varies as we increase R, and $\varepsilon_\tau/\varepsilon^*_{i,t}$ captures the limit as the return goes to infinity. To summarize

Proposition 1 *For hand-to-mouth individuals, the discount factor between period τ and t can be approximated by the probability ratio of the early winning probability over the switching point*

Figure 1: Bias vs headline prize when income increases

(a) Job seeker income levels (see Section 6): low (current) income = 668, high (future) income = 1440

(b) Income levels from Prolific validation experiment (see Section 5): low (current) income = 3567, high (future) income = 4048

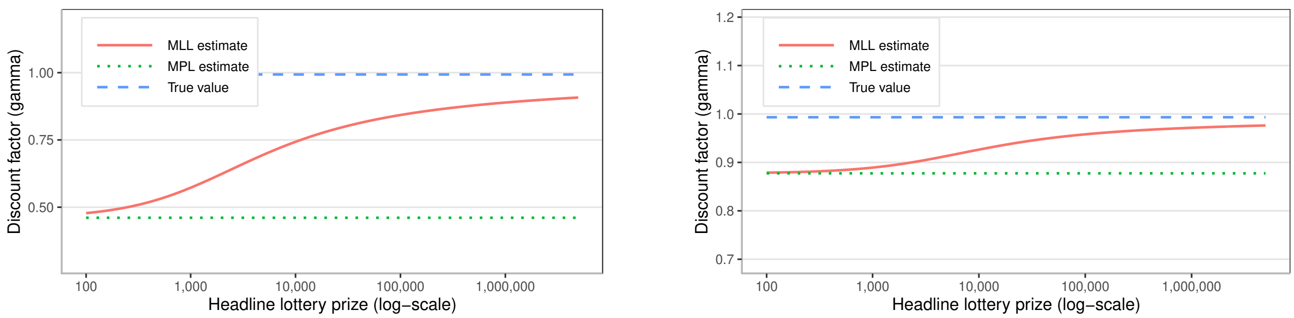

of the late winning probability when the reward becomes sufficiently large, independent of beliefs about other income.

The simple calibration that showed that a multiple price list approach would yield a biased discount factor estimate can be reproduced using a Multiple Lottery List (MLL) approach. Using the same parameters, we show the MLL estimates for increasing values of the headline lottery prize (R) in Figure 1, with job seeker income levels in panel (a) and income levels from our validation experiment in panel (b). For small prizes the bias is similar to the MPL estimate, but for a prize exceeding 10,000 dollar it is already much closer to the true value. For prizes above one million, the MLL estimate is close to the true value. Note that the MPL estimate shown in Figure 1 reflects a constant experimental reward, and not the case of an increasing reward as we discussed in footnote 9.[10] Furthermore, we show various alternative parameterizations of the calibration in Online Appendix B and always confirm the finding that MLL performs well with large headline prizes.

Testable implications: One obvious testable implication of our theory is that for large enough rewards MLL should not be affected by changes in background consumption while standard monetary methods such as MPL are. One can make these predictions sharper for particular types of changes. Consider a given individual with expectations $Eu_i(y_{i,t})$ at different times t_1, t_2 and t_3. Now measure the discount factor between the earlier periods ($\gamma_{i,t_1,t_2,M}$) and the later periods ($\gamma_{i,t_2,t_3,M}$) with a given elicitation method M. As an example, imagine an individual who always expects the high net income in Figure 1 and therefore any method measures the true discount factor according to the dashed blue line. Now assume that the same individual

[10]In Online Appendix B we show that even with increasingly large rewards, MPL fails to identify the true discount factor, in line with the limitations discussed in Online Appendix A.3.

realises that his unavoidable costs in the first period (C_{i,t_1}) are higher than initially thought, so that his new beliefs about his net income $\tilde{E}u_i(y_{i,t_1})$ are first order stochastically dominated by his initial beliefs $Eu_i(y_{i,t_1})$ while the other beliefs are not affected (where variables with "tilde" refer to those under the new beliefs). Again, measure the discount factors $\tilde{\gamma}_{i,t_1,t_2,M}$ and $\tilde{\gamma}_{i,t_2,t_3,M}$ under the new beliefs. Returning to the example, imagine that the individual now expects the low income in Figure 1 in the early period and the high income thereafter, and therefore the initial discount factor is measured according to the solid red line (under MLL) or dotted green line (under MPL). Proposition 1 implies that for sufficiently large lottery rewards MLL should not be affected by such changes in background consumption, while under standard monetary methods such as MPL the larger marginal utility of consumption in the early period is measured as a lower discount factor:

$$\tilde{\gamma}_{i,t_1,t_2,MLL} \approx \gamma_{i,t_1,t_2,MLL}$$
$$\tilde{\gamma}_{i,t_1,t_2,MPL} < \gamma_{i,t_1,t_2,MPL}.$$

Since there are no changes to the later periods, the measured discount factors there are not affected ($\tilde{\gamma}_{i,t_2,t_3,MLL} = \gamma_{i,t_2,t_3,MLL}$ and $\tilde{\gamma}_{i,t_2,t_3,MPL} = \gamma_{i,t_2,t_3,MPL}$). In the example this is satisfied as the red solid line is much closer to the true value of the blue line than the green dotted line. If individuals are quasi-hyperbolic discounters and $t_1 = 0$, this induces the following implications:

$$\tilde{\beta}_{i,MLL} \approx \beta_{i,MLL} \tag{5}$$
$$\tilde{\beta}_{i,MPL} < \beta_{i,MPL} \tag{6}$$

This is immediate from the previous comparison when $t_2 - t_1 = t_3 - t_2$ since one can directly infer the present-bias parameter as $\beta_{i,M} = \gamma_{i,t_1,t_2,M}/\gamma_{i,t_2,t_3,M}$. These are the predictions within an individual that experiences changes in beliefs. Similarly, if there are two groups of people and only those in group $\tilde{A}$ are aware of the higher costs while those in group A are not (i.e., the beliefs about net income of each individual in group $\tilde{A}$ is first order stochastically dominated by the beliefs of his/her counterpart in group A), then theory predicts (5) and (6) between these groups. Moreover, while we use MPL to illustrate these effects, other existing monetary elicitation methods share this feature when they do not take changing background consumption into account.[11] Readers interested in the empirical investigation of these predictions are invited

[11]For example, one also uncovers relationship (6) if one allows individuals not only to either take small reward r_τ early or r_t late, but additionally allows them to choose $\alpha \in [0, 1]$ so that they receive reward αr_τ early and $(1 - \alpha)r_t$ late as in the method of Convex Time Budgets of Andreoni and Sprenger (2012) when assuming fixed background consumption. The reason is simply that lower immediate consumption is associated with a higher marginal utility of immediate consumption which acts similarly to more present-bias.

to jump to Section 5 where we present a validation experiment.

To construct our new elicitation method in a way that distinguishes it from previous methods, we exploited two innovations: 1) shifting choices from the reward domain to the probability domain, and 2) making the stakes in some states of the world (e.g., after a lottery win) very large. The second is the key feature that enables us to undertake a limit argument. The first feature has several roles: it allows a simple trade-off between the likelihood of becoming rich and the timing; it makes this approach feasible as high prizes of the order mentioned in the previous paragraph would not be feasible unless they are paid out with small probabilities; and it reduces individual's incentives to borrow against future rewards already today. The last point might be a worry in methods where it is assumed that individuals are hand-to-mouth (i.e., no borrowing or saving) but individuals already know today that they will receive a large reward in the future which creates incentives to try to front-load their consumption. But if individuals do not know ahead of time if they will win and if the probability of winning is sufficiently small to render the expected return of the lottery low, then the incentives for ex-ante borrowing are low. We return to the last point more formally in the subsequent extensions (Section 4.2).

4 Important Theoretical Extensions

This section explores important theoretical extensions of our basic model, which we kept simple for illustrative purposes. In particular, our simple analysis so far assumed that an individual who wins the lottery early cannot spread the reward over multiple periods, either mechanically or through optimal savings. Our simple analysis also did not allow individuals to save or borrow against their normal income, the introduction of which usually implies that one measures their interest rate on savings rather than their time preferences and has been a major criticism to the use of monetary elicitation methods. We have also so far abstracted from other adjustments such as altering one's income process through job search. Our implementation so far has also been rather abstract, and here we discuss a more intuitive one via commercial lottery tickets. We assumed expected utility theory which relies on linearity in probabilities, and one might wonder how it would change if individuals weighted probabilities non-linearly in their utility. Finally, we assumed additively time-separable preferences so far, and study an alternative. These are discussed in turn.

The main insight is that the probability ratio which we introduced in the previous section continues to be unaffected by changes in normal income/consumption when the lottery win (R) tends to infinity. The probability ratio continues to uncover correctly which individual has a higher discount factor than another, and in many situations it continues to uncover the level of the discount factor correctly.

The treatment is not supposed to cover all complexities of borrowing/saving, departures from expected utility, or deviations from additive time-separable utility. Rather, it highlights that some of the main conclusions of our simple exercise do not immediately fall apart when the environment is extended into one direction or another. Since it is not possible to cover all conceivable theoretical extensions and their combinations here, we rely for broader concerns on the validation exercise in the subsequent section.

4.1 Spreading experimental payments over time: equal split or optimal savings after lottery win

Our baseline model that equates the discount factor with the probability ratio in (4) relies on hand-to-mouth consumers. While this has been a common assumption in the literature and might be palatable for regular incomes, alternatives abound.

Spreading rewards equally over time: Andersen et al. (2008) assumes that experimental payments are divided equally across a finite number N of subsequent periods. One can think of this akin to an annuity where rewards are spread evenly over time. It can also be viewed more agnostically as a way to introduce some finite planning horizon that individuals have in their mind when evaluating rewards. Applied to our setting, this means that a lottery win in period t leads to consumption $y_{i,t'} + R/N$ in periods $t' \in \{t, t+1, ..., t+N-1\}$, while consumption remains at $y_{i,t'}$ in all other periods. For long-lived individuals that live for more than $t+N$ periods the analogue to (4) is now adjusted for the fact that the difference between winning the lottery and not winning the lottery is spread over multiple periods:

$$\begin{aligned}
\gamma_{i,\tau,t} &= \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \frac{\sum_{n=0}^{N-1} \gamma_{i,\tau,\tau+n}[Eu_i(R/N + y_{i,\tau+n}) - Eu_i(y_{i,\tau+n})]}{\sum_{n=0}^{N-1} \gamma_{i,t,t+n}[Eu_i(R/N + y_{i,t+n}) - Eu_i(y_{i,t+n})]} &\quad (7)\\
&\approx \frac{\varepsilon_0}{\varepsilon^*_{i,t}} \frac{\sum_{n=0}^{N-1} \gamma_{i,\tau,\tau+n}}{\sum_{n=0}^{N-1} \gamma_{i,t,t+n}} \text{ for R large,} &\quad (8)
\end{aligned}$$

where the approximation follows the same logic as earlier. This can be simplified for quasi-hyperbolic discounters (who have $\gamma_{i,\tau,t} = \delta_i^{t-\tau}$, except that $\gamma_{i,0,t} = \beta_i \delta_i^t$ with $\beta_i \leq 1$). For them expression (8) immediately implies

$$\gamma_{i,0,t} \approx \begin{cases} \varepsilon_\tau/\varepsilon^*_{i,t} & \text{for } \tau > 0 \text{ and R large} \\ \varepsilon_\tau/\varepsilon^*_{i,t} \left(\frac{1+\beta_i \sum_{n=1}^{N-1} \delta_i^n}{1+\sum_{n=1}^{N-1} \delta_i^n}\right) & \text{for } \tau = 0 \text{ and R large} \end{cases} \quad (9)$$

That is, the probability-ratio still identifies the discount factor approximately if payments are all in the future. This follows directly since in this case $\gamma_{i,\tau,\tau+n} = \gamma_{i,t,t+n}$. If early payments are in the present ($\tau = 0$) the probability ratio $\varepsilon_0/\varepsilon^*_{i,t}$ overstates the true discount factor. Nevertheless, conditional on the long-run discount factor which can be correctly identified via (9), the probability ratio $\varepsilon_0/\varepsilon^*_{i,t}$ ranks individuals correctly in terms of their short-term discount factor. This is particularly useful as in many settings the goal is to identify those with higher or lower present-bias and correlate this with other traits, this remains possible.[12] We therefore obtain the following for long-lived individuals:

Proposition 2 *If rewards are equally spread over N periods but individuals are otherwise hand-to-mouth, their discount factor continues to be related to the probability ratio via*

$$\gamma_{i,\tau,t} \approx \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \tag{10}$$

as in (4) for sufficiently large lottery rewards if (i) the individual is an exponential discounter, (ii) the individual is a quasi-hyperbolic discounter and $\tau > 0$.

For quasi-hyperbolic discounters and immediate rewards ($\tau = 0$) the following holds in the limit when R grows large: the probability ratio overstates the actual short-run discount factor, but if individuals only differ in discount factors $\gamma_{i,0,t}$ the probability ratio ranks them correctly: for any two such individuals i and j it holds that $\gamma_{i,0,t} > \gamma_{j,0,t}$ if and only if $\varepsilon_\tau/\varepsilon^*_{i,t} > \varepsilon_\tau/\varepsilon^*_{j,t}$.

Savings after lottery win: Instead of equally spreading the win over a fixed number N of periods, one can assume that individuals can optimally save their rewards across these periods. This is substantially more involved, but we will see that the insights from Proposition 2 carry over to this setting. The rest of this subsection provides the details and is only relevant to the interested reader. While optimal savings can also be allowed across finite saving horizons N, to keep exposition tractable we will focus on the limit case where N is infinity for this exposition.

Consider therefore the same setting as before where individuals cannot save or borrow in normal times (which will be relaxed in the next subsection) but after winning the lottery assume that they have sufficient income to access capital markets of the simplest form: individuals can save at interest rate ι after a lottery win as long as their net worth remains positive, i.e., wealth plus net present value of future earnings remains above zero.

For this extension, we focus on exponential or quasi-hyperbolic discounters. Assume $\iota \leq \delta_i$ which is sufficient to ensure finiteness of their value functions. We focus on a population with homogeneous sophistication. Either all individuals are naive and expect their future selves

[12] The long-run discount factor can be obtained from, e.g., by eliciting $\varepsilon_1/\varepsilon^*_{i,2}$, which by (8) equals δ_i. Given δ_i and for high enough R, a higher value of β is associated with a higher value of the probability ratio since (9) can be rewritten as $\varepsilon_0/\varepsilon^*_{i,t} = \beta_i\delta_i^n[1 + \sum_{n=1}^{N-1}\delta_i^n]/[1 + \beta\sum_{n=1}^{N-1}\delta_i^n]$, which is increasing in β_i.

to be exponential discounters, so they expect their current optimal consumption plan to be implemented by future selves. Or all individuals are sophisticates and anticipate that future selves are also quasi-hyperbolic discounters, and they play a game with their future selves. We obviously allow regular income $y_{i,t}$ to vary across time periods t but abstract from uncertainty in income except for our lottery, which saves on notation and creates an environment where the savings game among sophisticates is well understood under the additional assumption that their Bernoulli utility has constant relative risk aversion above unity (Laibson (1996)).[13] We make this assumption on Bernoulli utility for sophisticates in the following. The equilibrium for finitely-lived individuals is unique, and we focus on the limit of these when we study infinitely lived individuals.[14] Online Appendix A.4 clarifies the setup further and discusses this point.

For the analysis, denote by $V_{i,\tau}$ the continuation value of person i from period τ onward when she does not win the lottery and consumes stream $\{y_{i,\tau}, y_{i,\tau+1},...\}$. Denote by $U_{i,\tau}(R + W_{i,\tau})$ the continuation utility of person i who wins the lottery in period t, has available a net present value of current and future income of $W_{i,\tau}$ in addition to the lottery win R, and can save. Note that this net present value lies in bounded set $[\underline{y}, \bar{y}/\iota]$. The utility of individual i at time zero who has the chance of receiving lottery win R with probability ε_τ in period τ is given by

$$(1 - \varepsilon_\tau)V_{i,0} + \varepsilon_\tau(V_{i,0} + \gamma_{i,0,\tau}(U_{i,\tau}(R + W_{i,\tau}) - V_{i,\tau})).$$

The first summand captures the utility when she does not win. The second captures the utility when she does win: She eats her income during the first $\tau - 1$ periods yielding utility $V_{i,0} - \gamma_{i,0,\tau} V_{i,\tau}$, i.e., the lifetime utility minus the continuation utility after the lottery win. After the lottery win she instead obtains continuation utility $U_{i,\tau}(R + W_{i,\tau})$, again discounted to the present.

The individual compares this expression with the analogous expression where the lottery happens at time t with probability ε_t. The point of indifference is given at a late winning

[13] For naive quasi-hyperbolic individuals we do not make additional assumptions on the Bernoulli utility, and for them it is trivial to allow for uncertainty in income, albeit at the cost of more notation.

[14] Sophisticates play a savings game with their future selves, which can allow additional equilibria (see Online Appendix A.4). This does not happen for naives.

probability $\varepsilon^*_{i,t}$ where both terms are equalized, which after rearrangement yields:[15]

$$\begin{aligned}
\gamma_{i,\tau,t} &= \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \frac{U_{i,\tau}(R+W_{i,\tau}) - V_{i,\tau}}{U_{i,t}(R+W_{i,t}) - V_{i,t}} \\
&\approx \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \frac{U_{i,\tau}(R+W_{i,\tau})}{U_{i,t}(R+W_{i,t})} \text{ for R large} \\
&\approx \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \frac{U_{i,\tau}(R)}{U_{i,t}(R)} \text{ for } R \text{ large.} \qquad (11)
\end{aligned}$$

The approximation involves two arguments. First, note that for large rewards the terms $V_{i,\tau}$ and $V_{i,t}$ drop out: They are both bounded because the income process is bounded and returns are discounted more than the interest rate; while $U_{i,t}$ and $U_{i,\tau}$ go to infinity as R goes to infinity. That continuation utilities after the lottery win do go to infinity follows trivially from the fact that at the time of the lottery win the individual could choose to consume all the proceeds and obtain a utility that approaches infinity, and any savings decision has to yield weakly higher utility. Second, note that the terms $W_{i,\tau}$ and $W_{i,t}$ drop out. This follows from the fact that they are bounded and the marginal utility $u'(c)$ of consumption tends to zero at very large consumption levels, and so when R is large additional units have small impacts. This is trivial in all finite horizon problems where consumption in all periods tends to infinity as the lottery reward rises, and marginal utilities converge to zero.[16] It is also well-known for standard cake-eating problems with infinitely-lived exponential discounters[17], which carry over to naive quasi-hyperbolics who expect to be exponential after the initial period. For infinitely-lived sophisticated quasi-hyperbolics with isoelastic utility and linear equilibrium savings rate Laibson (1996), equation 29, provides an explicit expression for $U_i(R)$, and it is easy to verify that $U_i'(R)$ goes to zero as lottery rewards approach infinity.

The important take-away is that the exact nature of future income sequences does not affect the measurement. The following proposition explores the consequences of this insight for infinitely-lived individuals in analogy to Proposition 2:

[15]The indifference condition is

$$\begin{aligned}
& (1-\varepsilon_\tau)V_{i,0} + \varepsilon_\tau(V_{i,0} + \gamma_{i,0,\tau}(U_{i,\tau}(R+W_{i,\tau}) - V_{i,\tau})) \\
= \; & (1-\varepsilon^*_{i,t})V_{i,0} + \varepsilon^*_{i,t}(V_{i,0} + \gamma_{i,0,t}(U_{i,t}(R+W_{i,t}) - V_{i,t}))
\end{aligned}$$

which reduces to the one in the text given that $\gamma_{i,\tau,t} = \gamma_{i,0,t}/\gamma_{i,0,\tau}$.

[16]While this might be less obvious in the game between selves of a sophisticated quasi-hyperbolic consumer, Laibson (1996) shows that they consume a constant fraction of wealth in each period, and it is easy to see that the marginal utility of that strategy goes to zero for additional income.

[17]For exponential discounters, standard arguments establish that initial period's consumption c in the savings problem is related to next period's wealth R_+ and consumption c_+ via $u'(c) = \delta_i U_i'(R_+)$. Since initial consumption c goes to infinity as initial resources R go to infinity, marginal utility of wealth has to go to zero. Since the same holds for naive quasi-hyperbolics from the second period onward, a similar logic applies to them.

Proposition 3 *In the setting with infinitely-lived individuals who can save after a lottery win, their discount factor continues to be related to the probability ratio via*

$$\gamma_{i,\tau,t} \approx \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \tag{12}$$

as in (4) for sufficiently large lottery rewards if (i) the individual is an exponential discounter, (ii) the individual is a sophisticated or naive quasi-hyperbolic discounter and $\tau > 0$.

For quasi-hyperbolic discounters and immediate rewards ($\tau = 0$) the following holds in the limit when R grows large: the probability ratio overstates the actual short-run discount factor, but if individuals only differ in discount factors $\gamma_{i,0,t}$ the probability ratio ranks them correctly: for any two such individuals i and j it holds that $\gamma_{i,0,t} > \gamma_{j,0,t}$ if and only if $\varepsilon_\tau/\varepsilon^*_{i,t} > \varepsilon_\tau/\varepsilon^*_{j,t}$.

Proof. *See Online Appendix A.5.* ■

In essence, with an infinite horizon condition (11) converts to (12) because the period in which the lottery win is received does not affect the continuation utility. That means that the probability ratio exactly identifies the long-run discount factor independently of whether or not individuals have options to save, how present-biased they are, or what their future earnings expectations are. For quasi-hyperbolics this does not hold in the first period where the future is more heavily discounted leading to a ratio $U_{i,t}(R)/U_{i,\tau}(R) > 1$. Still, this effect does not overturn the direct effect of higher impatience, all else equal, and the probability ratio rises in β_i. So we can rank individuals correctly according to their present-bias. This ranking compares individuals with identical long-run discount factor δ_i, but by (12) this discount factor can be elicited correctly and one can control for it. Therefore, even in the short-run the probability ratio serves as a sensible device to rank individuals.

In the Online Appendix B.2, we extend the simple calibration exercise from section 3 to the case where savings are allowed after a lottery win. We demonstrate that also in this case our lottery-based elicitation method is able to identify discount factors fairly accurately despite income fluctuations. We show that this finding persists for both naïve and sophisticated individuals and is robust against different parameters values (such as income levels or the true discount factor).

4.2 General savings opportunities do not influence discount factor measurement

It has been argued that standard discount factor elicitation experiments do not just measure time preferences, but rather capture the degree of credit constraints - see for example Dean and Sautmann (2020) or the review in Cohen et al. (2020). The reason is simple: Those who can freely borrow and save at some interest rate should accept choices with higher net present value evaluated at their interest rate, and then use savings and borrowing to transfer that net present value across time in the desired way. [18] Choices then reflect their interest rate, and not their time preferences. We will see that this is not the case in our elicitation method.

Consider a similar setting as in the previous subsection, but allow individuals to save both before and after the lottery win. Focus on naive quasi-hyperbolic discounters who can save/borrow at a person-specific interest rate ι_i in normal times, and they can do so at rate ι_i^L in any period after winning our lottery. We distinguish the time after winning since individuals might hire financial advisors after winning a high-stakes lottery. Assume interest rates are below individuals' discount rate to avoid discussions of unbounded solutions for infinitely-lived consumers. Standard transversality conditions apply.

Note that these individuals can borrow and save, but lack sophisticated insurance instruments and can only use savings to adjust the expected return. Fix R at a very high level so that the results in the previous section provide a reasonable approximation, but then reduce the winning probability ϵ_τ towards zero so the expected return stays low. This seems a reasonable approach given that the lotteries we propose have a market value of only a few Euro. In this case our method recovers the true discount factor and not the interest rate. In analogue to Proposition 3 we obtain:

Proposition 4 *Consider infinitely-lived naive quasi-hyperbolic discounters who can save at person-specific interest rate τ_i^L in any period after winning our lottery and at rate τ_i otherwise. Fix two individuals. For R sufficiently large, and then ϵ_τ sufficiently small, the results from Proposition 2 continue to apply: the probability ratio approximates the true discount factor for each of the individuals when the time of the early lottery is in the future ($\tau > 0$); and when the early lottery is in the present the probability ratio ranks the individuals correctly according to their short-run discount factor if the individuals are identical except for present-bias and person-specific interest rate τ_i.*

Proof. See Online Appendix A.6. ■

The key insight here is the standard logic from existing envelope theorems: consider the consumption sequence $c_{i,1}, c_{i,2}, c_{i3,}...$ of individual i in the absence of our elicitation method. Now introduce our elicitation method using a fixed very high reward R so that the approximation in our method is approximately valid, but a low winning probability so that the expected returns are minor. The choices for naive quasi-hyperbolic individuals or exponential discounters are characterized by a standard inter-temporal optimization problem, and we can rely on the

[18] This need not be an issue if individual's bracket narrowly, as pointed out by Andreoni et al. (2018).

envelope theorem to see that changes in their regular consumption (in absence of a lottery win) are of second-order impact. When these individuals evaluate which lottery to choose, their utility will essentially be determined by the timing and probability of winning, but not by minor adjustments in their normal consumption. We can therefore treat these individuals as if their background consumption $c_{i,1}, c_{i,2}, c_{i3,...}$ is fixed.

Nothing in the argument up to this point is specific to our elicitation method. It applies equally to standard elicitation methods. Yet, those usually assume that the sequence of background consumption is constant, whereas individuals that can save and borrow would not choose such a constant consumption sequence except in the very special case where the time discount rate equals the interest rate. The strength of our method is that it is designed to uncover the discount factor independently of the exact sequence of background consumption $c_{i,1}, c_{i,2}, c_{i3,...}$, and this is what allows us to apply it as outlined in the previous subsection.

While we consider here only individuals who can borrow and save, it is possible to include individuals who have more severe borrowing constraints. For example, consider individuals who cannot save at all except after a lottery win. Their probability ratio still uncovers their discount factor between future periods and ranks them correctly even against the savers in terms of short-run discount factor, as the exact nature of consumption during normal times does not matter as shown in the previous section. Our proof does not, however, cover sophisticated individuals who can borrow and save even in the absence of the lottery win. For them the envelope theorem does not apply, as their future consumption is determined by a game against their future selves rather than a classical optimization problem. Whether a different proof technique can establish a similar result for such individuals is beyond the scope of this paper.

4.3 Implementation via commercial lottery tickets

In our baseline theory we consider the probability ratio $\epsilon_\tau/\epsilon^*_{i,t}$ and show that this can be equated to the discount factor if the reward conditional on winning tends to infinity as, e.g., in equation (4).

A much more intuitive and practical implementation is via commercial lottery tickets, each of which pays out a reward R with probability $\hat{\epsilon}$. Consider at most N lottery tickets, and assume that the consumer is indifferent between n_τ lottery tickets at time time τ, or $n^*_{i,t}$ lottery tickets at a later time t. A naive interpretation would be to assume that therefore the discount factor must be roughly equal to the ratio of the number of early lottery tickets over late lottery tickets, i.e.,

$$\gamma_{i,\tau,t} \approx n_\tau/n^*_{i,t}. \tag{13}$$

This would be supported by the naive interpretation that $\epsilon_\tau = n_\tau\hat{\epsilon}$ and $\epsilon^*_{i,t} = n^*_{i,t}\hat{\epsilon}$. Obviously this naive approach abstracts from the possibility of winning on multiple lottery tickets at the same time. Still, this intuition is approximately valid when we first consider R large, but then let $\hat{\epsilon}$ be sufficiently small to keep the expected reward constant at a low level (which in our applications is below 2 Euro). This is the case for commercial lottery tickets that have a small purchase price, and offer high rewards at very small probabilities. We make this point precise in the Online Appendix A.7.

4.4 Non-expected utility

Our basic model in Section 3 relies on expected utility theory. A particularly popular alternative relies on probability weighting: the individual does not consider probabilities p directly, but through a strictly increasing weighting function $w(p)$ with domain $[0, 1]$ (see, e.g., the overview in Starmer (2000)). To illustrate its effect, consider preferences according to rank-dependent expected utility (see, e.g., Quiggin (1982), Machina (1994), or cumulative prospect theory going back to Tversky and Kahneman (1992)). With only two outcomes, the probability of the more attractive one gets evaluated by $w(p)$, while the less attractive one get evaluated by one minus this, so that the sum still equals unity. To keep this notational simplicity we focus again on a time-varying but deterministic income stream $y_{i,t}$ so that uncertainty enters through our lottery, though this can easily be extended to stochastic income streams.

The key concept to adjust the previous discount factor formula depends on the elasticity of the weighting function. Let ϕ denote this elasticity evaluated at $p = 0$. Consider $\phi \in (0, \infty)$, as is the case of the parametric forms used in Tversky and Kahneman (1992), Masatlioglu and Raymond (2016), and others. Focus again on the case of a fixed expected return from the lottery in the early period (i.e., $\varepsilon_\tau R$ is fixed) and let R become large. In Online Appendix A.8 we show that we now obtain instead of (4) the adjusted formula

$$\gamma_{i,\tau,t} \approx \frac{w(\varepsilon_\tau)}{w(\varepsilon^*_{i,t})} \approx 1 - \phi + \phi \frac{\varepsilon_\tau}{\varepsilon^*_{i,t}}, \tag{14}$$

where the first applies under our standard approximation of R large. The second is valid only if discounting is not too strong so that $\varepsilon_\tau / \varepsilon^*_{i,t}$ is close to unity. It has the advantage that it nicely showcases a key component that is important under probability weighting, namely the elasticity of the weighting function at low probabilities which governs the slope of the affine transformation that maps the probability ratio into the discount factor. While the probability ratio now does not directly uncover the discount factor, either approximation shows clearly that individuals with higher revealed probability ratio have a higher discount factor.[19] Therefore, the

[19] This is obvious in the affine case, and the first approximation can be written as $\gamma_{i,\tau,t} \approx \frac{w(\varepsilon_\tau)}{w(\varepsilon_\tau/\zeta)}$ where ζ

probability ratio ranks individuals in the same way as their discount factor.

Note that this discussion relied on all individuals sharing the same probability weighting function. That is, the probability weighting function was not indexed by identifier i. Interestingly, in the parametric versions of Tversky and Kahneman (1992) and Masatlioglu and Raymond (2016), even if individuals are heterogeneous with respect to the parameters of their weighting function, the elasticity at $p = 0$ is identical across them. So in these cases heterogeneity of the probability weighting function does not lead to heterogeneity in ϕ, and therefore allows consistent ranking of individual discount factors by the probability ratio. In Masatlioglu and Raymond (2016) it even holds that $\phi = 1$, in which case the original formula $\gamma_{i,\tau,t} \approx \varepsilon_\tau/\varepsilon^*_{i,t}$ applies even in the presence of probability weighting with heterogeneous weighting parameters.

Prospect theory often encompasses not only probability weighting but also theories of reference dependence. If individuals have an exogeneous reference point, a sufficiently large reward conditional on winning clearly lies in the gain domain, and our derivations have shown that the exact payoffs conditional on not winning the lottery do not matter as long as they are bounded (and not affected by the lottery win). In this case it can easily be shown that such reference dependence does not further affect our results. In some formulations of reference dependence the reference point is endogenously determined by expectations about future outcomes. Then the impact of our lotteries are less obvious, as these directly affect the reference point and, thereby, the utility especially after not winning the lottery. In rational expectation based models of disappointment aversion such as Bell (1985) and Loomes and Sugden (1986), the reference point is the expected utility of the future outcome. To understand the impact of our lotteries on this reference point, assume that the expected monetary payout from our lotteries is bounded and that individual utility functions have vanishing marginal utility, as for example with CRRA utility. In this case we can show that with rising reward R conditional on winning the impact of the lottery on the reference point converges to zero.[20] The reason is that marginal utility is not rising fast enough in the reward to compensate for the vanishing winning probability. So for large rewards it is as if the reference point is exogeneously given. Under the assumptions of this section the result extends also to the choice-acclimating personal equilibrium of Kőszegi and Rabin (2007).[21] If marginal utility is not vanishing, the impact of the lottery on the reference

equals the probability ratio $\varepsilon_\tau/\varepsilon^*_{i,t}$. Therefore, a strictly increasing weighting function implies that a higher ζ is associated with a higher discount factor.

[20] The rational expectations reference point for a person contemplating any lottery in period t is given by $(1-\epsilon_t)Eu_i(y_{i,\tau}) + \epsilon_\tau Eu_i(y_{i,t} + R)$ which is bounded below by the reference point $Eu_i(y_{i,\tau})$ in absence of the lottery. Under the assumption that the product of the winning probability times the reward is bounded ($\epsilon_t R < B$ for some B), an upper bound for the reference point is $(1-B/R)Eu_i(y_{i,\tau}) + E[u_i(y_{i,t}+R)B/R]$ which converges to its lower bound since $u_i(y_{i,t}+R)B/R \to 0$ when $R \to \infty$ as long as $u'(R) \to 0$, by L'Hopital's rule.

[21] In general the theory of Kőszegi and Rabin (2007) is more complex that that of Bell (1985) and Loomes and Sugden (1986) because it does not feature a single-valued reference point against which all outcomes are evaluated, but each outcome is separately compared to each outcome in a reference lottery. Under the specific

point vanishes only if the expected monetary payoff becomes exceedingly low.[22] This shows a variety of settings where reference dependence does not further affect our analysis.

4.5 Separating intra-temporal risk aversion from inter-temporal elasticity of substitution in consumption

Most of the literature on the experimental elicitation of discount factors uses the additively time-separable utility function that underlies our baseline specification in Chapter 3. Nevertheless, it has been pointed out that such specifications restrict atemporal risk aversion to coincide with the intertemporal elasticity of substitution in consumption (e.g., Epstein and Zin (1989)), and tractable alternatives that separate the two concepts have been proposed.

In the context of discount factor elicitation, Andersen et al. (2018) highlight that additive time-separability coupled with hand-to-mouth consumption (i.e., no savings) also implies indifference to inter-temporal correlation of consumption, which they reject in empirical work using stochastic rewards. In their theory and in their estimation they use a particularly tractable extension of the baseline framework, which we adopt here as well for our purposes to provide a basic insight into non-time-separable utility. They assume that consumption c_i at two periods in time τ and t is evaluated according to

$$E\left[\varphi\left(\gamma_{i,0,\tau}u_i(c_{i,\tau}) + \gamma_{i,0,t}u_i(c_{i,t})\right)\right]. \tag{15}$$

where all notation is identical as in our baseline model in Chapter 3, except for function $\varphi(x) = x^{1-\eta}/1-\eta$ with $\eta \in (0,1)$.[23]

They assume income is directly consumed, as in our baseline model. Allowing for time-variation in background income and associated consumption but abstracting from uncertainty,

assumption of this section that background consumption can be individual and time-dependent but has no risk, and given that our lottery has only two possible outcomes, we are in the binary risk case in a given period and for this one both approaches are identical (see O'Donoghue and Sprenger (2018) for a comparison of these approaches).

[22] Assume that marginal utility does not vanish (i.e., $u'(R) \to D > 0$) and the expected monetary payment of the lottery is bounded by B ($\epsilon_t R < B$ for some B), then for sufficiently large R the reference point $(1-\epsilon_t)Eu_i(y_{i,\tau}) + \epsilon_\tau Eu_i(y_{i,t}+R)$ is bounded above for example by $(1-B/R)Eu_i(y_{i,\tau}) + 2DRB/R]$. Given such an R, the reference point becomes therefore becomes arbitrarily close to $Eu_i(y_{i,\tau})$ for B sufficiently small, where $Eu_i(y_{i,\tau})$ constitutes the reference point of in the absence of a lottery.

[23] Andersen et al. (2018) use utility function $u_i(c) = c^{1-r}/1-r$ and estimate $r \in (0,1)$, with $r \approx 0.3$, which fits our assumptions: increasing, concave and unbounded. They estimate $\eta \approx 0.3$, so within set $(0,1)$ that we use here for exposition. They highlight that an analogue to (15) can be derived as in Epstein and Zin (1989) when abstracting from the timing of the resolution of uncertainty, i.e., resolving it in the same period.

we obtain as utility for our lottery in the early period τ:

$$\frac{(1-\varepsilon_\tau)\left(\gamma_{i,0,\tau}u_i(y_{i,\tau})+\gamma_{i,0,t}u_i(y_{i,t})\right)^{1-\eta}}{1-\eta}+\frac{\varepsilon_\tau\left(\gamma_{i,0,\tau}u_i(y_{i,\tau}+R)+\gamma_{i,0,t}u_i(y_{i,t})\right)^{1-\eta}}{1-\eta}$$

We obtain an analogous expression for the lottery in the later period t:

$$\frac{(1-\varepsilon_t)\left(\gamma_{i,0,\tau}u_i(y_{i,\tau})+\gamma_{i,0,t}u_i(y_{i,t})\right)^{1-\eta}}{1-\eta}+\frac{\varepsilon_t\left(\gamma_{i,0,\tau}u_i(y_{i,\tau})+\gamma_{i,0,t}u_i(y_{i,t}+R)\right)^{1-\eta}}{1-\eta}$$

Let $\varepsilon^*_{i,t}$ denote the late lottery that renders individual i indifferent between both. Setting both expressions equal, using $\gamma_{i,\tau,t}=\gamma_{i,0,t}/\gamma_{i,0,\tau}$ and re-arranging yields:

$$\frac{\varepsilon_\tau}{\varepsilon^*_{i,t}} \quad = \quad \frac{\left(u_i(y_{i,\tau})+\gamma_{i,\tau,t}u_i(y_{i,t}+R)\right)^{1-\eta}-\left(u_i(y_{i,\tau})+\gamma_{i,\tau,t}u_i(y_{i,t})\right)^{1-\eta}}{\left(u_i(y_{i,\tau}+R)+\gamma_{i,\tau,t}u_i(y_{i,t})\right)^{1-\eta}-\left(u_i(y_{i,\tau})+\gamma_{i,\tau,t}u_i(y_{i,t})\right)^{1-\eta}}$$

Obviously, for $\eta=1$ specification (15) becomes again additively time-separable, and we obtain again exactly the condition we exploited in our baseline in Section 3. In general, we loose the time-separability, but at a very large R we still obtain approximation:

$$\frac{\varepsilon_\tau}{\varepsilon^*_{i,t}}\approx\frac{\left(u_i(y_{i,\tau})+\gamma_{i,\tau,t}u_i(y_{i,t}+R)\right)^{1-\eta}}{\left(u_i(y_{i,\tau}+R)+\gamma_{i,\tau,t}u_i(y_{i,t})\right)^{1-\eta}}\approx\frac{\gamma^{1-\eta}_{i,\tau,t}u_i(y_{i,t}+R)^{1-\eta}}{u_i(y_{i,\tau}+R)^{1-\eta}}\approx\gamma^{1-\eta}_{i,\tau,t},$$

where the first and second approximation follow because the omitted terms remain finite, and the last approximation follows by the same logic that established (4) in our baseline model.

Measurement is again independent of time-variation in background consumption, and a higher probability ratio on the left identifies a higher discount factor on the right. We do not have to elicit parameter η to rank individuals with respect to their discount factor, as long as individuals share this parameter that governs how they deviate from additive time-separability (note that they can continue to differ in utility $u_i(.)$ and in background consumption $y_{i,\tau}$ and $y_{i,t}$). This can be investigated, and we take some comfort in two observations: First, Andersen et al. (2018) investigate heterogeneity of η by observables and report only two out of fifteen coefficients to be significant (indicating that one might want to control for age and gender). Second, in our subsequent validation exercise our method correlates well with an established elicitation method when there is no variation to background consumption, even though our method introduces intertemporal risk while the other does not. If our method performs well when background consumption is constant, this section confirms our central message that variation over time in background consumption does not alter the relationship between the probability ratio and the discount factor.

4.6 A general model with job search, with or without savings.

Finally, our analysis can be embedded into a richer overall model. The introduction of this paper and motivation for this work is based around settings where incomes might vary substantially in short amounts of time, which includes for example environments where unemployed individuals search for jobs. DellaVigna and Paserman (2005) outline the consequences of hyperbolic discounting in an extension of the sequential search paper of McCall (1970), but they did not directly measure discount factors. Their model can be recast as follows: an infinitely lived individual obtains utility b when unemployed and makes two choices per period: She can exert job search effort a at contemporaneous cost $k(a)$, and she can choose a reservation utility level r that the offer at least has to generate to get accepted. The job search effort generates probability a of receiving a job offer this period, and a job offer specifies a wage w during each period of employment in this job where employment starts the following period. The offer is a random draw from distribution $G(.)$ with non-zero density on $[w_L, w_H]$. The job seeker decides to accept and offer or not. Employment continues with probability $1-q$ from one period to the next, and terminates with probability q leading to unemployment.

In the Additional Material we present a larger model that nests this setup, with or without savings, and allows costs $k(\cdot)$ and offers $G(\cdot)$ to be individual-specific and have time subscripts that can capture changes with age or unemployment duration. Time subscripts are also relevant if the costs $k_t(a)$ represent a setting where the returns to search effort depend on changing aggregate conditions through a larger matching-function in the Diamond-Mortensen-Pissarides tradition.[24] Propositions 1 - 4 remain valid even in this extension. The reason for this is that in the derivation of these propositions, the utility conditional on not winning the lottery becomes negligible relative to the utility after winning the lottery, which implies a trade-off that is driven by the timing and probability of winning rather than the details of what happens in the absence of a lottery win. We return to measuring discount factors and correlating them with job search outcomes in the spirit of DellaVigna and Paserman (2005) in Section 6.

[24] In a standard macro-search model such as Pissarides (2000) the number of matches $M(AU,V)$ in a given period depends on the current measure of unemployed individuals (U) multiplied by their current average search effort (A) and the current number of vacancies (V). An individual takes these as given, and her own matching probability depends on her search effort $\tilde{a}$ according to $\tilde{a}M(AU,V)/(AU)$. If the cost of effort is $\tilde{k}(\tilde{a})$, then we can rewrite this equivalently as $a := \tilde{a}M(AU,V)/(AU)$ and $k(a) = \tilde{k}(aAU/M(AU,V))$. If A, U or V are changing over time, this implies the derived k to change over time even if $\tilde{k}(\cdot)$ and $M(\cdot,\cdot)$ are not time dependent. We will therefore allow for time subscripts in our extended model.

5 Validation experiment

There might be concerns about our lottery method that exceed those explored in the previous theoretical extensions: among other things individuals may feel differently about lotteries leading to answers that might shed little light on "usual" time preferences; or they might have particular behavioral biases in low-probability high-return situations that we have not modelled. Rather than exploring all these theoretically, we perform a validation exercise. We compare our method to a well-established alternative ('convex time budgets' (CTB), see Andreoni and Sprenger (2012)) and explore whether randomly induced variations in expected future expenditures (through "prompting", explained below) affect elicited preferences in either method. We perform both a between subject and within subject comparison. We find that with relatively stable expenditures, the two methods produce preference estimates that are significantly correlated. When we induce fluctuations in expected future expenditures, the estimates from our method remain much more stable than those elicited using CTB.

5.1 Experimental setup

We validated our method experimentally using an online study involving 305 participants from the research platform Prolific. The experiment was conducted in July 2023. Participants were required to reside in New York State and have an annual income of at most $99,999[25]. The study took place on July 10th 2023.[26]

Participants were asked to complete 4 consecutive blocks of questions: (i) a block of 20 convex-budget-set questions (CTB, Andreoni and Sprenger (2012)), (ii) a block of 20 multiple lottery list-questions (MLL), (iii) one question on risk aversion ('the bomb-question' proposed by Crosetto and Filippin (2013)) and (iv) a number of questions on demographics, income and expenditure. The precise order was subject to treatment assignment and is explained below.

The CTB and MLL questions have equivalent rates of return ranging from 20% to 100%. Early options are either the week of the experiment or 2 months after it, and late options are either 2 months or 4 months after the early option. CTB questions consist of an early and a late option that vary in time of payout and token-exchange-rate; each participant has 100 tokens per question to allocate between the two options. MLL questions consist of an early and a late option that vary in the likelihood of receiving a lottery-scratch-ticket and timing; each participant chooses for each question either the early or late option. One question out of all CTB

[25] The median household income in 2022 was $75,157 and the Average Household Income was $111,583 (Source: US Census data)

[26] The study was preregistered at the AEA RCT Registry (Belot et al. (2023)). The pre-analysis plan (available as supplementary material 4) states an intended sample size of 400, but due to a lower number of interested participants the ultimate sample is 312.

and MLL questions is randomly selected for payout. For instructions and complete questions see the Supplementary Material 2. The selected parameters for the blocks of questions for CTB and MLL are summarized in Table C.1 in the Online Appendix.

Upon completion of the study participants received a participation fee of $10 plus additional earnings that depended on their response to the randomly selected question. The participation fee and potential earnings from CTB questions were automatically administered through their Prolific account. If a MLL question was randomly selected for payout, a second random draw would determine whether they earned the lottery ticket (with the odds according to their response).

5.2 Lottery tickets

The lottery tickets were New York State lottery tickets, with a headline prize of $5 million.[27] If a participant earned a lottery ticket, we scratched the ticket on their behalf and informed them of the outcome on the date corresponding to their choice. If the ticket won the headline prize, the participant would be invited to collect the winning ticket in person.[28]

These lottery tickets combined with our procedure had various advantages. First, it allowed us to offer participants a lottery with a single and very large headline prize, in line with our theory. Second, we could determine the date on which participants learn whether they won or not (as opposed to lottery tickets with a pre-defined drawing date). Third, obtaining the tickets did not require any actions from the participant (such as creating an online account) nor did we have to physically mail tickets to participants. Both ensured that participants did not face any additional hurdles or uncertainty for obtaining their ticket.

As a result of our setup, the transaction costs were kept constant *within* method across choices at different time horizons: CTB monetary payoffs were paid out through Prolific, lottery ticket outcomes were sent through a message on the participant's Prolific account. There was a slight delay (up to three days) for immediate payouts, but this was the same for both methods and clearly communicated to participants.

5.3 Prompting treatment and randomization

Our theory predicts that preferences elicited with CTB and MLL are similar when future expenditures are constant, but different when future expenditures fluctuate. A validation exercise thus

[27]The tickets also offer smaller prices with larger odds, but participants were informed that they were only eligible to receive the headline prize. Smaller prizes were donated to charity.

[28]The scratching of tickets was recorded and available to watch by the participants to guarantee reliability of the process. Participants were informed about the precise procedures at the start of the study and provided consent to the possibility of receiving a lottery ticket. The Pre-Analysis plan provides all the relevant details

requires random variation in whether individuals perceive their future expenditures as stable or not. Clearly such variation is difficult to generate. We use an approach inspired by insights from psychology, which shows that saliency increases the weight of what is made salient into the decision (see Taylor and Thompson (1982)). The intervention we propose includes a series of detailed questions about expenditures coming up in the coming month. By bringing these immediate expenditures to the forefront of people's mind, we would expect that they would then behave as if they are facing higher immediate expenditures than those who have not been prompted to think about them. There is recent evidence in the context of a developing country (Zambia) that prompting people to think about their expenditures does indeed seem to affect their perceived expenditures and financial decisions (Augenblick et al. (2022)). If this intervention succeeds as intended, it should affect decisions in a similar manner as if the treated individuals truly experienced a variation in their immediate expenditures. The idea is that treated individuals realize that their short-term expenditures are higher than expected, inducing them to feel poor in the short-term and to internalize a lower short-term background consumption.

We designed the experiment to allow for both a between and within subject comparison. Table 1 describes the four treatments. In Treatment 1 and 2 we first elicit preferences using both CTB and MLL (in different order). Subsequently these groups answer the prompting questions before answering one set of elicitation questions once more. Each group only repeats one elicitation method (CTB in group 1 and MLL in group 2) to avoid lack of attention from an excessive number of questions. From these groups we can test whether responses change after prompting. This approach has largest statistical power because it makes within-individual comparisons. However one may worry that, due to 'anchoring', individuals are less likely to reconsider their responses when asked the same set of questions a second time. This concern may attenuate the potential impact of the prompting intervention. Therefore we also implement treatments 3 and 4, in which participants first answer the prompting questions and then the elicitation questions (both CTB and MLL). A comparison of responses from groups 3 and 4 with the first responses of groups 1 and 2 provides a valid between-subject design. While statistical power is lower for this test due to heterogeneity across individuals, this approach does not suffer from the potential anchoring bias. All treatment groups complete the survey by answering a question on risk preferences and a set of questions on demographics that also includes expectations about expenditures and income in the coming six months. Formal hypotheses (including a test of the first-stage) are provided below.

Table 1: Experimental design

Treatment 1	Treatment 2	Treatment 3	Treatment 4
		Prompting	**Prompting**
CTB	MLL	CTB	MLL
MLL	CTB	MLL	CTB
Prompting	**Prompting**		
CTB	MLL		
Risk	Risk	Risk	Risk
Demographics	Demographics	Demographics	Demographics

5.4 First stage and balance

Individuals were randomly assigned to one of the four treatments. As expected, the treatment groups are balanced in terms of the demographic characteristics that we collected (see Table C.2 in the Online Appendix). The core hypothesis of the experiment relies on the experimental intervention to successfully induce a change in perceived current expenditures. We first check whether the intervention had the intended effect. Specifically, our "first-stage hypothesis" is that the 'prompted' participants report a higher estimate of their total immediate expenditures than the non-prompted participants.

To check that our experimental intervention did indeed induce a perception of larger estimates of expenditures in the short run, the survey includes four sets of questions. First, a general question (asked to all respondents) about total expected expenditures in the upcoming month (Expenditure_{pre}). Second, the individuals were asked to report expected expenditures for a detailed list of categories (see supplementary material for all categories). Third, a computed total from the categories was shown to the respondents and they were asked to confirm this total number or provide a different number if not correct ($\text{Expenditure}_{post}$). Finally, at the end of the experiment all respondents were asked to report their predicted monthly expenditures and income for each of the next 6 months.

As a test of the first-stage we compare expenditure estimates *before prompting* in Treatment groups 1 and 2, with the expenditure estimates *after prompting* in Treatment group 3 and 4. Our hypothesis is:

First-stage hyp. 0.1: $E(\text{Expenditure}_{\text{pre}}|\text{Treat} = 1, 2) \ < \ E(\text{Expenditure}_{\text{post}}|\text{Treat} = 3, 4)$

Similarly, we compare within-individual responses before and after the prompting questions:

First-stage hyp. 0.2: $E(\text{Expenditure}_{\text{post}} - \text{Expenditure}_{\text{pre}}|\text{Treat} = 1, 2) \ > \ 0$

Table 2: First-stage tests (hypotheses 0.1 and 0.2)

	Pre-prompting	Post-prompting	P-value one-sided t-test
Hyp. 0.1 (between)			
Expenditures, mean	2,519.61	3,269.57	0.002
Standard error mean	112.69	242.90	
Observations	152	145	
Hyp. 0.2 (within)			
Expenditures, mean	2,519.61	3,400.13	0.000
Standard error mean	112.69	195.40	
Observations	152	152	

Excluding respondents for whom the ratio of the pre-prompting response to the post-prompting response was above 3.

The tests of these hypotheses are presented in Table 2. In both cases we find that the intervention strongly increased expenditures estimates (with p-values < 0.01). The difference is large: around \$700 in the between comparison and around \$900 in the within comparison.[29] In appendix Figures C.1 and C.2 we provide further graphical evidence that confirms our first-stage hypotheses. These figures also show that the increase in perceived expenditures only relates to immediate expenditures: the post-prompting expenditure estimate for the immediate period ('the coming 3-4 weeks') is significantly larger than the post-prompting estimates for each of the subsequent three months.[30] That is, only immediate expenditures are perceived higher.[31]

We conclude that the prompting intervention had the intended effect and succeeded in generating an exogenous variation in perceived immediate expenditures.[32]

[29]We impose some data cleaning to deal with extreme outliers in estimated expenditure responses. In particular we exclude individuals for whom pre-prompting expenditure estimates exceed 3 times the post-prompting estimate that is based on categories (and explicitly confirmed). We show that alternative ways of dealing with outliers does not alter our conclusion that the intervention had the desired impact (see Tables C.3 and C.4 in the Online Appendix).

[30]This is tested using a within individuals t-test, leading to p-values of 0.01 or smaller in all three tests.

[31]Interestingly, this difference does not seem to be driven only by a saliency of 'unusual expenditures' since only 20 % of respondents report an unusual expenditure of EUR 400 or more. Rather, the difference appears to stem also from the detailed listing of regular expenditures

[32]Hypothesis 0.2 includes only treatment groups 1 and 2 in order to directly relate to our core hypothesis below. Instead one might include all four treatment groups, which yields the same conclusion: the post-prompting expenditure are 708 dollar higher (p-value < 0.001).

5.5 Core hypotheses

We consider quasi-hyperbolic discounters and define our core hypotheses in terms of the β (present-bias) parameters estimated under different methods. Our pre-registered hypothesis use empirical analogues to the predictions (5) and (6) in Section 3, where individuals after prompting realize that they need to spend less and correspond to those with higher perceived short-term cost. We compare our MLL method to CTB (rather than MPL) because it is more commonly used nowadays.

We estimate β as well as δ (long-run *monthly* discount factor) from the CTB and MLL question responses. For CTB we follow closely the approach outlined by Andreoni and Sprenger (2012). For MLL we take the midpoint of the range that is identified in the common multiple price list approach. The estimation details can be found in the Online Appendix C.1.[33]

The core hypothesis is that variations in short-term background consumption do not affect discount factor measurement under MLL. We do expect such variation to translate into significant differences in measured time preferences under the CTB method. We will test this both with a between-individual and a within-individual design. From the between-subject comparisons, we expect present bias to be stronger on average under CTB with prompting than under CTB without prompting:

Hyp. 1.1 The average $\hat{\beta}_{i,CTB}$ for individuals in treatments (3) and (4) is significantly lower than the average $\hat{\beta}_{i,CTB-NP}$ for individuals in treatments (1) and (2).

For MLL we expect no such difference:

Hyp. 1.2 There is no significant difference between the average $\hat{\beta}_{i,MLL}$ for individuals in treatments (3) and (4) compared to the average $\hat{\beta}_{i,MLL-NP}$ for individuals in treatments (1) and (2).

Within-individual we expect present bias to be stronger under CTB with prompting than without, i.e., the within-individual difference in measured β is significantly negative. Within-individual we expect to find no significant difference between β under MLL with and without prompting:

Hyp. 2.1: The difference $\hat{\beta}_{i,CTB-P} - \hat{\beta}_{i,CTB-NP}$ is significantly negative for individuals in treatment (1).

[33]Summary statistics for the estimated parameters are shown in the Online Appendix: Figure C.3 and Table C.5.

Table 3: Core hypotheses

	No prompting	Prompting	P-value one-sided t-test
Hyp. 1.1 (between)			
Beta (CTB estimates), mean	1.005	0.997	0.398
Standard error mean	0.020	0.019	
Observations	157	155	
Hyp. 1.2 (between)			
Beta (MLL estimates), mean	0.987	0.980	0.326
Standard error mean	0.012	0.011	
Observations	148	148	
Hyp. 2.1 (within)			
Beta (CTB estimates), mean	1.038	0.980	0.049
Standard error mean	0.028	0.025	
Observations	79	79	
Hyp. 2.2 (within)			
Beta (MLL estimates), mean	0.982	0.978	0.427
Standard error mean	0.020	0.021	
Observations	71	71	

Hypothesis 1.1 tests $\hat{\beta}_{CTB,NP} = \hat{\beta}_{CTB,P}$ between samples. Hypothesis 1.2 tests $\hat{\beta}_{MLL,NP} = \hat{\beta}_{MLL,P}$ between samples. Hypothesis 2.1 tests $\hat{\beta}_{CTB,NP} = \hat{\beta}_{CTB,P}$ within sample (paired test). Hypothesis 2.2 tests $\hat{\beta}_{MLL,NP} = \hat{\beta}_{MLL,P}$ within sample (paired test). All t-tests are one-sided. $\hat{\beta}$ in all cases censored to values in the range 0.5-1.5 (replacing more extreme values with these limits).

Hyp. 2.2: The difference $\hat{\beta}_{i,MLL-P} - \hat{\beta}_{i,MLL-NP}$ is not significantly different from zero for individuals in treatment (2).

Table 3 shows results for our four core hypotheses.[34] The between-individual comparison of CTB estimates of β with and without prompting (hyp. 1.1) does not yield a significant difference. Also the equivalent comparison for MLL estimates (hyp. 1.2) does not yield a significant difference. The within-individual comparisons do show a significant impact of expenditure variations (induced by prompting) on CTB estimates: we find significantly more present bias after prompting han before prompting ((0.980 vs 1.038, p-value 0.049). The same does not hold for the MLL estimates: after prompting the mean β does not change significantly. These results support our

[34]The CTB estimates for β and δ contain a number of outliers which have large impacts on the sample means. In our baseline results we censor all estimates outside the range 0.5-1.5. Alternatives are provided in the appendix (see Tables C.6 and C.7). Conclusions are similar, although statistical significance is diluted when too many outliers are included.

theory. When fluctuations in expenditures increase, preferences as estimated from CTB change, while MLL estimates remain similar. While we do not confirm this from the between-subject comparison, lack of statistical power might play a role here.[35]

5.6 Secondary hypotheses

The core hypotheses compare mean estimates of the present-bias parameter. As secondary hypotheses we also consider individual level correlations in estimates rather than only the mean.

For individuals that answer both MLL and CTB without prompting (i.e., individuals in treatments (1) and (2) for their first two sets of questions), we expect a significant positive correlation between $\hat{\delta}_{i,MLL-NP}$ and $\hat{\delta}_{i,CTB-NP}$. For these individuals we also expect a significant positive correlation between $\hat{\beta}_{i,MLL-NP}$ and $\hat{\beta}_{i,CTB-NP}$. This indicates that individuals who are ranked higher in terms of long-run discount factor or present-bias in one methodology also tend to be ranked higher under the other methodology in the absence of prompting. To test this hypothesis we will run a regression of the form:

$$\hat{\delta}_{i,MLL-NP} = \gamma_0 + \gamma_1 \hat{\delta}_{i,CTB-NP} + \varepsilon_i \tag{16}$$

Hyp. 3.1: A similar ranking of the MLL and CTB estimates in the absence of prompting would result in a significantly positive slope coefficient γ_1 in this regression. We perform a similar regression for present bias estimates β.

Hyp. 3.2: We expect a weaker correlation between $\hat{\beta}_{i,MLL}$ and $\hat{\beta}_{i,CTB}$ elicited under prompting in treatments (3) and (4). We might not see a correlation between them as MLL measurements should be unaffected by income variation while CTB measurements are affected by income variation. We investigate this hypothesis using a regression similar to (16).

Results are visualized using binned scatterplots in Figure 2. In panel (a) we find a strong positive correlation between CTB and MLL estimates for δ elicited without prompting, supporting hypothesis 3.1. The corresponding rank correlation is 0.54. Panel (b) shows that the association is substantially weaker under prompting (hypothesis 3.2). The same pattern emerges for β: the association between the CTB and MLL estimates is substantially stronger without prompting (panel (c)) than with prompting (panel (d)).

The corresponding regression coefficients are presented in Table C.8 in the Online Appendix

[35]It might not be surprising that the between-individual design did not detect the effects. If the effect size found in hypothesis 2.1 were true, our power to detect it in the between-individual comparison is only 66% due to the larger variance in discount factors across individuals.

Figure 2: Binned scatter plots comparing MLL and CTB estimates (hypotheses 3.1 and 3.2)

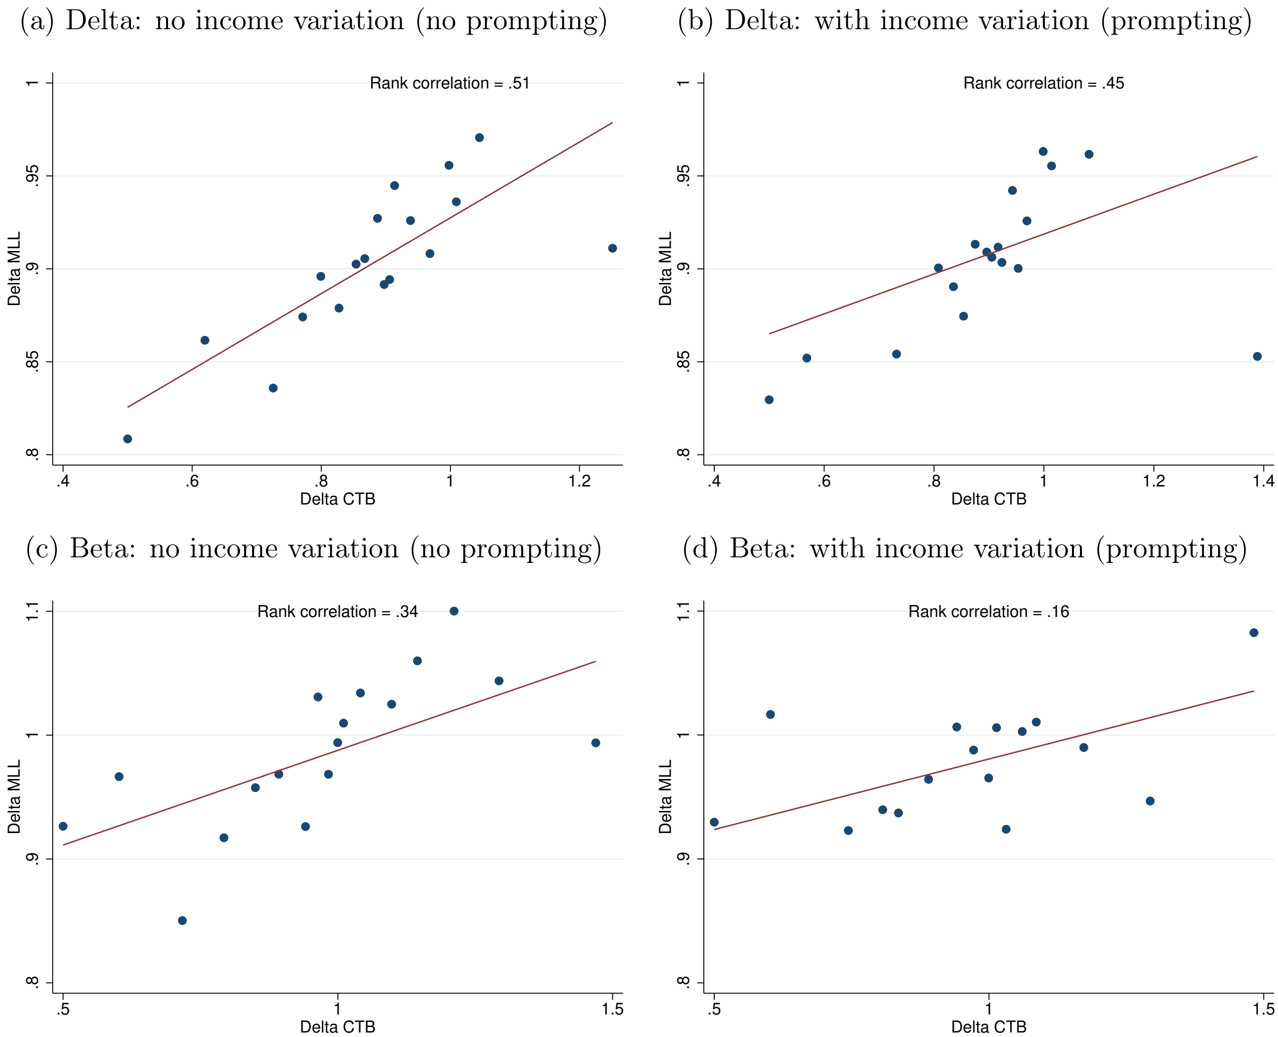

and confirm these results. We find significant slope coefficients in all cases, but they are larger in magnitude for the no-prompting samples than in the prompting samples, both for δ (0.204 vs 0.107, the difference being significant with p-value 0.024) and for β (0.153 vs 0.114, the difference not significant). These secondary findings are again in line with our theory: CTB and MLL estimates are more aligned when expenditures are stable than when expenditures fluctuate.

To assess the association between the MLL and CTB estimates we might reverse the regression and have the CTB estimate on the left-hand-side.[36] Results are qualitatively similar, as presented in Table C.9, with slope coefficients even closer to unity for the samples without prompting.

[36]We thank an anonymous referee for this suggestion.

5.7 Exploratory analysis

As a final (exploratory) analysis we consider within-individual correlations between estimates before and after prompting (possible only for treatments (1) and (2)). We expect the correlation between $\hat{\beta}_{i,MLL-P}$ and $\hat{\beta}_{i,MLL-NP}$ for individuals in treatment (1) to be stronger than the correlation between $\hat{\beta}_{i,CTB-P}$ and $\hat{\beta}_{i,CTB-NP}$ for individuals in treatment (2). Figure 3 shows that this is precisely the case. For MLL (panel (a)), there is a strong positive relation (rank correlation = 0.54) between the pre- and post-prompting estimates. The analogous rank correlation for CTB is only 0.18 as (see panel (b)) and not statistically significant, signalling that responses changed significantly after our intervention highlighted high current expenditures. The corresponding regression coefficients confirm these findings (see Table C.10 in the Online Appendix), with a slope coefficient for MLL of 0.566 (significant at the 1% level) and 0.197 for CTB (not statistically significant).

Figure 3: Binned scatter plots comparing within individual estimates before-after prompting (corresponding to the regressions in Table C.10)

(a) Beta: MLL estimates

(b) Beta: CTB estimates

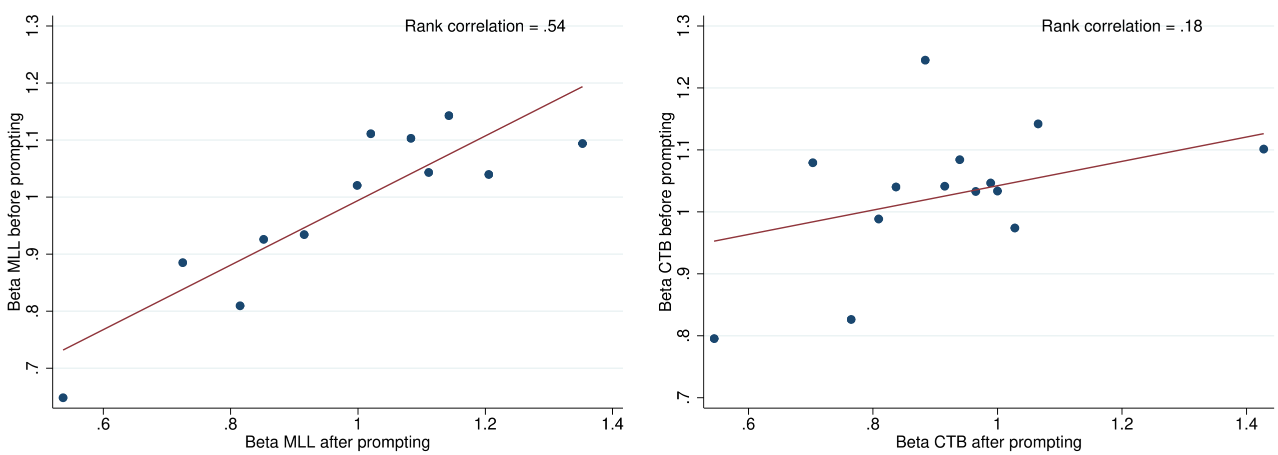

5.8 Conclusion

Concluding, we find three main results that support our theory. First, the mean MLL estimates for β are the same before and after prompting, while CTB mean estimates differ before and after prompting. Second, CTB and MLL estimates at the individual level correlate stronger without prompting than with prompting. And third, at an individual level the pre-prompting MLL estimates for β strongly correlate with the post-prompting estimates. This is not the case for CTB estimates.

Note that we also performed another experiment with a student population in 2019-2020, which has been reported in earlier drafts of the current paper. The design had important shortcomings, that we have overcome with the current design. For completeness we include the details and results of this other experiment as supplementary material.[37]

6 Application to Job Search

We now return to our initial motivation and apply our proposed methodology to examine the role of time preferences in job search. Such a connection has been stressed in DellaVigna and Paserman (2005), and more recent contributions argue both for its relevance (e.g., DellaVigna et al. (2017), DellaVigna et al. (2020)) and against (e.g., Marinescu and Skandalis (2020)). To our knowledge, there is only one recent study that attempts to correlate measures of time preferences with elicited measures of job search. Le Bihan (2023) elicits time preferences in a sample of French unemployed job seekers and correlate these measures with measures of job search effort and reservation wages. He uses 'conventional methods' (Double Multiple Price List and CTB), both in the monetary and effort domains. The contribution is also methodological as he shows that the different methods lead to different estimates of time preferences and correlate differently with measures of job search behavior.

Our goal here is to examine the question of how time preferences affect job search using a validated method for eliciting time preferences under varying background consumption. We elicited time preferences from job seekers during a field experiment with 300 unemployed job seekers in Scotland in 2013/2014 (see Belot et al. (2019b) for experimental details). These were measured with our MLL method, in a very similar way to those deployed in our validation experiment, this time with a headline prize of £250000. We refer to the Online Appendix D for details on how we elicited these preferences and calculated beta and delta parameters.

Figure D.1 in Online Appendix shows the distributions of time preference parameters in our sample. The mean and median of beta are centered at one, with little heterogeneity. That means, we find little evidence of present bias using our method. There is more heterogeneity is the long term discount factor, with the 25th to 75th percentile spreading between around 0.5 and 0.93.[38] We did not deploy other methods to elicit time preferences due to constraints on the number of questions we could include. For validation of our method see the previous section.

As part of the larger study, we collected measures of job search effort and reservation wages at baseline and over the 12 weeks panel dimension of the experiment (Belot et al. (2019a)). The

[37] We thank an anonymous referee for their suggestion to implement an improved elicitation experiment design.

[38] Lottery tickets were of comparable value, costing £2 in the UK. They had a headline prize of £250000 and also offered lower prizes.

study introduced an experimental variation at week 4, where half of the sample received suggestions of alternative occupations to include in their search. The effects of this intervention are evaluated in Belot et al. (2019b). To avoid differences arising from the experimental treatment, we use only data from the first three weeks of the study here.

Our sample is too small to examine job finding, which is a rare event in our setting, but we can use job interviews as a more frequent proxy for which our sample is sufficiently powered (see the discussion in Belot et al. (2019b) on this). We can also directly test the predictions by DellaVigna and Paserman (2005) regarding the relationship between time preference parameters with search effort and with reservation wages.

Relating time preferences to search effort and job interviews. We have two empirical measures of search effort: (1) self-reported hours spent searching in the previous week, and (2) number of applications. The first measure is obviously the most direct measure of effort. The other is plausibly positively correlated with search effort but the relationship is less clear. For example, people may spend a lot of hours searching but may be more selective and not necessarily send more applications. On the other hand, the first measure is perhaps more subject to reporting error.

For the reservation wage, we also have several alternative measures. The first is a self-reported measure collected at baseline where we ask individuals the minimum wage they would be willing to accept. The second is a measure based on the actual search behavior of individuals. Since a large share (over 40%) of vacancies include information about wages, we can examine what the lowest wage is that individuals consider in their search.[39]

Finally, we have one usable measure of job search success: the number of job interviews. This is often taken as a proxy for job finding. A caveat to this through the lens of job search models in the spirit of DellaVigna and Paserman (2005) is that even a successful job interview might not lead to job acceptance if the reservation wage is too high. This might be less of an issue in Scotland where a large share of firms indicate wage offers in their job postings, and therefore job seekers might be able to avoid interviews at jobs where the offers are too low.

DellaVigna and Paserman (2005) build on the model of McCall (1970) where individuals choose a reservation wage, and in their extension also a level of search effort. They embed quasi-hyperbolic discounting into this setting. They predict that less present bias (higher β) leads to more job search, because individuals do not procrastinate. In addition, it has no effect on reservation wages for naive agents, as reservation wages are decided today but only affect payments in the future. Taken together, lower present bias is predicted to increase job finding.

Table 4 presents the results. Point estimates for the effect of present bias are noisily estimated

[39]Note that the first outcome (stated reservation wage) is observed once per individual, while the definition based on search behavior is observed at most three times per individual.

Table 4: Regressions job search effort and log reservation wage

	(1) Weekly hours spent on job search	(2) Weekly applications	(3) Min. (hourly min. wage)	(4) Log hourly stated res. wage	(5) Weekly interviews
Beta	-1.645	0.040	0.020	0.075	-0.268
	(2.860)	(0.308)	(0.023)	(0.133)	(0.282)
Delta	-5.127*	0.384	-0.014	-0.030	-0.805***
	(3.035)	(0.542)	(0.029)	(0.153)	(0.098)
Model	Linear RE	Poisson RE	Linear RE	Linear	Poisson RE
N	500	692	655	205	642

Standard errors (clustered by individual) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All models contain control variables (risk preferences, age, gender, white, couple, children, higher education, unemployment duration) . All models except (4) contain individual random effects. Presented regression results are from a Poisson regression for interviews and number of applications. We observe these outcomes, as well as the reservation wage (col (3)) during each of the first 3 weeks of the study and thus have (at most) three observations per individual. To account for unobserved heterogeneity we include individual random effects.

and insignificant. To the extent that point estimates still give some indication: we see that for hours spent searching (column 1) and the likelihood of a job interview (column 5), the point estimates are negative, while DellaVigna and Paserman (2005) predict that those with low present-bias (high β) search more and are more likely to find a job. For the reservation wage, the results are shown in columns 3 and 4. Effects of present bias on either measure of the reservation wage are economically small and insignificant, in line with DellaVigna and Paserman (2005).

Regarding the impact of the long-run discount factor, the model of DellaVigna and Paserman (2005) predicts that individuals with higher long-run discount factor (higher δ) provide more search effort as they have more interest in its future rewards, and have higher reservation wages as they are more willing to wait for better future offers. The first increases job finding, the second decreases it, so the overall effect cannot be signed. The second effect dominates when the long-run discount factor is close to one.

Table 4 reveals that those with a higher long-run discount factor have significantly fewer job interviews. They have significantly lower hours of job search, while we find no relationship with job applications. There is no economic or statistically relevant relationship with either measure of reservation wages. These findings are not as predicted by DellaVigna and Paserman (2005).

Discussion of our findings on job search and time preferences. Our results do not provide much evidence of present bias in job seekers, and as a consequence, we do not find much impact on job search behavior. The findings on the long-run discount factor are more significant, though not in line with the predictions from existing work such as DellaVigna and Paserman (2005). We briefly discuss what might be missing.

First, reservation wages are a difficult concept: Krueger and Mueller (2016) report that 44% of job seekers in their study accepted a job below their stated reservation wage, and Hall and Mueller (2018) back out a much larger variation in non-wage characteristics than in wages. A reservation utility might then manifest itself more on non-wage characteristics. This could be explored in future work.

Second, if those with a higher long-run discount factor target only jobs with more attractive non-work characteristics, they might not need to search over as big a set of jobs to send their applications. To the extent that this can be pre-specified through, e.g., tighter search radius or other filters, this reduces the set of alternatives that the individual has to consider. Even if the effort to inspect each alternative remains constant, this would reduce the overall time searching. Even if the number of applications remained the same, applying to jobs with better characteristics might lower their interview and job finding chances. A form of partially-directed search along these lines might reconcile the findings above.

Both of these could in principle be investigated with our method combined with data from modern job search platforms, though this might require a larger sample of job seekers.

7 Conclusion

This paper lays out a simple method for eliciting time preferences even in settings where income and consumption might change over time. It relies on high stakes lotteries, where the main trade-off is between the timing and the probability of getting rich. If the amount of money associated with winning the lottery is sufficiently high, level and variation in the normal income and consumption stream do not affect this trade-off. Our method identifies the discount factor in simple settings, and preserves the ranking and sometimes even the level in more elaborate settings with savings or probability weights. In an experimental setting, we find that the discount factors elicited with this method indeed seem to be unaffected by shocks. An illustration with job seekers showcases an application in a setting where income variation is the norm.

If appropriate lottery tickets are available, our method is relatively simple and fast to implement. It might therefore serve in other settings where income fluctuations are a worry. Future work might explore more deeply how our method performs relative to, say, structural estimation techniques that incorporate elicitations of subjective expectations about future variation in income and consumption. We discussed the pros and cons of either approach already in the introduction, and view this as an open agenda.

References

Ambrus, A., Ásgeirsdóttir, T. L., Noor, J., and Sándor, L. (2014). Compensated Discount Functions: An Experiment on the Influence of Expected Income on Time Preferences. Working Paper 197736, Harvard University OpenScholar.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2008). Eliciting risk and time preferences. *Econometrica*, 76(3):583–618.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2013). Discounting behaviour and the magnitude effect: Evidence from a field experiment in denmark. *Economica*, 80:670–697.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2018). Multiattribute utility theory, intertemporal utility, and correlation aversion. *International Economic Review*, 59(2):537–555.

Andreoni, J., Gravert, C., Kuhn, M. A., Saccardo, S., and Yang, Y. (2018). Arbitrage or narrow bracketing? on using money to measure intertemporal preferences. Nber working paper no 25232.

Andreoni, J. and Sprenger, C. (2012). Estimating time preferences from convex budgets. *American Economic Review*, 102(7):3333–56.

Attema, A. E., Bleichrodt, H., Gao, Y., Huang, Z., and Wakker, P. P. (2016). Measuring discounting without measuring utility. *American Economic Review*, 106(6):1476–94.

Augenblick, N., Jack, K., Kaur, S., Masiye, F., and Swanson, N. (2022). Budget neglect in consumption smoothing: A field experiment on seasonal hunger. Technical report, Working paper.

Augenblick, N., Niederle, M., and Sprenger, C. (2015). Working over Time: Dynamic Inconsistency in Real Effort Tasks *. *The Quarterly Journal of Economics*, 130(3):1067–1115.

Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations research*, 33(1):1–27.

Belot, M., Kircher, P., and Muller, P. (2019a). Data for: Providing advice to jobseekers at low cost: An experimental study on online advice. *The Review of Economic Studies*, 86(4):1411–1447. https://doi.org/10.1093/restud/rdy059.

Belot, M., Kircher, P., and Muller, P. (2019b). Providing advice to jobseekers at low cost: An experimental study on online advice. *The Review of Economic Studies*, 86(4):1411–1447.

Belot, M., Kircher, P., and Muller, P. (2023). Time preferences using a multiple lottery list: validation. *AEA RCT Registry, AEARCTR-0011729, https://doi.org/10.1257/rct.11729-1.1*.

Belot, M., Kircher, P., and Muller, P. (2024). Data and code for: Eliciting time preferences when income and consumption vary: Theory, validation & application to job search. Technical report, American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. http://doi.org/10.3886/E198281V1.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

Caliendo, M., Cobb-Clark, D. A., and Uhlendorff, A. (2015). Locus of control and job search strategies. *Review of Economics and Statistics*, 97(1):88–103.

Carvalho, L. S., Meier, S., and Wang, S. W. (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review*, 106(2):260–84.

Cohen, J., Ericson, K. M., Laibson, D., and White, J. M. (2020). Measuring time preferences. *Journal of Economic Literature*, 58(2):299–347.

Coller, M. and Williams, M. B. (1999). Eliciting individual discount rates. *Experimental Economics*, 2(2):107–127.

Crosetto, P. and Filippin, A. (2013). The "bomb" risk elicitation task. *Journal of Risk and Uncertainty*, 47(1):31–65.

Dean, M. and Sautmann, A. (2020). Credit constraints and the measurement of time preferences. *The Review of Economics and Statistics*, 0(ja):1–52.

DellaVigna, S., Heining, J., Schmieder, J. F., and Trenkle, S. (2020). Evidence on job search models from a survey of unemployed workers in germany. Working paper.

DellaVigna, S., Lindner, A., Reizer, B., and Schmieder, J. F. (2017). Reference-dependent job search: Evidence from Hungary. *The Quarterly Journal of Economics*, 132(4):1969–2018.

DellaVigna, S. and Paserman, M. D. (2005). Job search and impatience. *Journal of Labor Economics*, 23(3):527–588.

Dohmen, T. (2014). Behavioral labor economics: Advances and future directions. *Labour Economics*, 30:71–85. Special Section articles on "What determined the dynamics of labour economics research in the past 25 years? edited by Joop Hartog and and European Association of Labour Economists 25th Annual Conference, Turin, Italy, 19-21 September 2013 Edited by Michele Pellizzari.

Dynarski, S. and Gruber, J. (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity*, 28(1):229–303.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937.

Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424.

Gerber, A. and Rohde, K. I. (2015). Eliciting discount functions when baseline consumption changes over time. *Journal of Economic Behavior and Organization*, 116:56–64.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica*, 89(5):2303–2339.

Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica*, 83(1):335–352.

Hall, R. I. and Mueller, A. I. (2018). Wage dispersion and search behavior: The importance of nonwage job values. *Journal of Political Economy*, 126(4):1594–1637.

Harrison, G. W., Lau, M. I., and Williams, M. B. (2002). Estimating individual discount rates in denmark: A field experiment. *American Economic Review*, 92(5):1606–1617.

Holden, S. T., Shiferaw, B., and Wik, M. (1998). Poverty, market imperfections and time preferences: of relevance for environmental policy? *Environment and Development Economics*, 3(1):105–130.

Kőszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073.

Krueger, A. B. and Mueller, A. I. (2016). A contribution to the empirics of reservation wages. *American Economic Journal: Economic Policy*, 8(1):142–179.

Laibson, D. I. (1996). Hyperbolic discount functions, undersaving, and savings policy. Working Paper 5635, National Bureau of Economic Research.

Laury, S. K., McInnes, M. M., and Swarthout, J. T. (2012). Avoiding the curves: Direct elicitation of time preferences. *Journal of Risk and Uncertainty*, 44(3):181–217.

Le Bihan, M. (2023). *Unemployment and job search: a behavioral and experimental approach.* PhD thesis, Université Lumière-Lyon II.

Loomes, G. and Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, 53(2):271–282.

Machina, M. J. (1994). Review of 'Generalized expected utility theory: The rank-dependent model'. *Journal of Economic Literature*, 32(3):1237–1238.

Marinescu, I. and Skandalis, D. (2020). Unemployment Insurance and Job Search Behavior. *The Quarterly Journal of Economics.*

Masatlioglu, Y. and Raymond, C. (2016). A behavioral analysis of stochastic reference dependence. *American Economic Review*, 106(9):2760–82.

McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 84:113–126.

Meier, S. and Sprenger, C. (2010). Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, 2(1):193–210.

Meyer, C. J. (2018). In search of a better life: Self-control in the ethiopian labor market.

Mueller, A. I., Spinnewijn, J., and Topa, G. (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review*, 111(1):324–63.

O'Donoghue, T. and Sprenger, C. (2018). Reference-dependent preferences. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 1, pages 1–77. Elsevier.

Pender, J. L. (1996). Discount rates and credit markets: Theory and evidence from rural india. *Journal of Development Economics*, 50(2):257–296.

Pissarides, C. A. (2000). *Equilibrium unemployment theory.* MIT press.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13(1):130–167.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, 38(2):332–382.

Taylor, S. E. and Thompson, S. C. (1982). Stalking the elusive" vividness" effect. *Psychological review*, 89(2):155.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.