

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

Michele Belot, Philipp Kircher, and Paul Muller*

April 13, 2021

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 lottery tickets with high rewards. In a standard intertemporal choice model high rewards decouple lottery choices from variation in background consumption. We validate our elicitation method experimentally on two student samples: one asked in December when their current budget is reduced by extraordinary expenditures for Christmas gifts; the other asked in February when no such extra constraints exist. We illustrate an application of our method with unemployed job seekers which naturally have income/consumption variation.

Keywords: time preferences, experimental elicitation, job search, hyperbolic discounting

JEL-codes: D90, J64

**Affiliations:* Belot, Cornell University and University of Edinburgh; Kircher, Cornell University, Catholic University of Louvain and University of Edinburgh; Muller, VU University Amsterdam and Tinbergen Institute. This study was pre-registered on the AEA RCT Registry, AEARCTR-0005115. We thank the Behavioral Lab of the School of Business and Economics (Vrije Universiteit Amsterdam) for data collection support, and many seminar audiences for their valuable suggestions. We particularly thank Ted O'Donoghue and Andrew Clausen. Kircher gratefully acknowledges support through European Research Council Grant No 284119.

1 Introduction

Imagine a researcher who is interested in the link between discounting and job search (as emphasized, e.g., in DellaVigna and Paserman (2005), DellaVigna et al. (2017)). Imagine she wants to investigate this by eliciting the short-run and long-run discount factor of individual job seekers, and then correlate these with their job search effort and success. To do this, she has to overcome the following hurdle: experimental elicitation methods based on monetary rewards evaluate the marginal utility of money at different points in time, which is affected both by their discount factor but also by their expectations of how much money they will have available for consumption in different periods. Concretely, unemployed individuals tend to have limited income currently but more income once they find a job, and 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)).

In such a setting, a job seeker who expects to find a job quickly without much effort would prefer money today because she has less need in the future, while an otherwise identical job seeker who expects that she has to search hard and long to find a job is more willing to take money in the future as she also expects to be poor then.¹ Without adjusting for the changing income stream, 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 other settings where individuals face uninsured risk because of income or expenditure shocks (such as medical or other extraordinary expenditures) raise similar concerns.

This issue has received little attention so far.² The current elicitation

¹The 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.

²A recent exception is Dean and Sautmann (2020) who point out that the measures of

methods such as multiple price lists (Coller and Williams (1999), Harrison et al. (2002)), convex budget sets (Andreoni and Sprenger (2012)), 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 changes in background consumption, or would require eliciting precisely how income and expenditures will change over the time. Even those methods where utility terms cancel and time preferences are directly recovered require a constant background consumption, otherwise the cancellation argument does not apply.

The researcher has several options to deal with this. 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)), the difficulty of eliciting a complex set of expectations might be the reason why this approach has not been extended to time-varying background consumption. Alternatively, 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, while a person who expects to be unemployed for longer might prefer to wait, even if their true discount factors are identical.

In this paper, we extend the idea of direct methods that aim to elicit time preferences without estimating utility functions to settings where income and associated background consumption change over time. We suggest a simple

time preferences cannot be interpreted as such in the presence of income shocks but can be used to inform about credit constraints individuals face.

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

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 the logic along two dimensions: probability weighting and savings.

To explore this idea, we rely on an elicitation method as in 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. In the extensions it is not always the level of the discount factor that is uncovered, but sometimes only 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. Our method is closest to Laury et al. (2012) except that they consider moderate rewards conditional on winning, while we consider very large rewards. As we will show, the size of the rewards is crucial to be able to back out the discount factors when background consumption varies over time, as it allows us to cancel utility terms in the limit. In their work, they assume background consumption does not change over time.

Our baseline model follows much of the literature in assuming that in-

dividuals consume any income they receive immediately. This assumption has raised two concerns in other work - one concerning savings after getting an experimental payment and one concerning savings before participants get experimental payments. After a payment, participants might spread it over multiple periods (see, e.g., Andersen et al. (2008) where individuals are allowed equal spreading over a fixed number of periods). We investigate theoretically how our method performs when individuals spread their experimental win over time through savings. For infinitely-lived hyperbolic discounters our simple measure still correctly identifies their long-run discount factor, without the need to estimate a structural model. With some qualifications our measure also ranks individuals correctly according to their short-run discount factor, though it overstates its magnitude.

The second criticism concerns heterogeneity in savings rates prior to an experimental payment: if intertemporal smoothing is possible, individuals should choose the option with the highest net present value discounted at their interest rate, and then use borrowing/saving to distribute it optimally over time. So their choices reflect their return on savings rather than their intrinsic discount factor (see, e.g., Cohen et al. (2020), Dean and Sautmann (2020)). We conclude the theory section by noting that this is not the case in our setting, at least for naive hyperbolic discounters who can save/borrow but lack complex financial instruments to insure against our lottery itself. At time zero they can only smooth the (small) expected win of the lottery as they are not yet told about success or failure. This ensures that our method still uncovers their true long-run discount factor. As long as saving opportunities after the lottery win are identical for example due to the ability to hire financial advisors, it also ranks otherwise similar individuals correctly according to their short-run discount factor even if their normal saving rates differ.

While real effort elicitation methods offer an existing avenue to avoid the problem of savings, they remain prone 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.

The fact that we exploit limit arguments involving large rewards at low probabilities begs the question whether our method still picks up the discount factors that we associate with usual rewards. We can explore this in settings where background consumption does not vary, where we can compare our measure with conventional elicitation methods. Ideally, we expect our measure to capture similar variation when background consumption is stable over time; and to differ in setting where it varies.

We therefore validate our method experimentally on two (student) samples recruited from the same underlying pool but asked to make choices at different times: one group is asked in December at a time when their current budget is reduced by extraordinary expenditures for Christmas and Saint Nicholas gifts while their budget in the near future is lifted by receipt of gifts. The other one is asked in February where we have no reason to expect difference in budget between current and future periods. Students only participated either in December or February. They are subjected to a standard time preference elicitation method via convex budget sets (Andreoni and Sprenger (2012)), as well as to our method.

Our main hypothesis is as follows: For the February group without the "Christmas shock", we expect both measures to behave similarly: both should correlate significantly in a regression, in theory with a coefficient of 1. In the December session - where shocks are expected in the near future - the correlation should be weaker as one measure was designed to be robust against such shocks while the other was not. We find clear confirmation of this hypothesis: regressing our measure on the measure according to the convex budget set method yields a highly significant positive coefficient in the February sample. The coefficient is no longer significant in the December sample. Still, conditional on self-reported absence of forthcoming shock, the long-run discount factors of both measures correlate significantly, with a coefficient that is close to one. Levels of discount factors between methods are not identical, but they pick up the same relevant variation in "normal" times but not under shocks.

Finally, we illustrate how our method can be deployed in the context of our leading example: job search. Since unemployed individuals naturally ex-

perience varying income streams associated with re-employment, we initially developed our method as a quick way to elicit discount factors despite this obstacle during a field experiment with 300 unemployed job seekers (see Belot et al. (2019b) for experimental details).⁴ 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 insignificant) 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 are not aware of previous 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. 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 and reports lower job search and job finding for those with higher measures of present bias.

The rest of the paper is structured as follows. Section 2 presents the theory in the simplest model by building on classical methods and expanding to ours, and Section 3 explores robustness to probability-weighting and savings. Section 4 presents the validation with student samples, and Section 5 presents our discount factor measures and correlations with job search activities in a sample of job seekers. Section 6 concludes.

⁴This is the first time we report on this part of the experiment. It preceded the full development of all parts of the theory presented here, and the validation exercise.

2 Time preference measurement: A Simple Model

Consider an individual i who cannot save or borrow, and faces exogenous and potentially stochastic income stream $y_{i,t}$ in periods $t = 0, 1, \dots, T$, where T could be infinite. Assume possible income realizations are in a bounded subset $[\underline{y}, \bar{y}]$ of the positive reals. 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, strictly concave Bernoulli utility function on the domain of the non-negative reals. 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 that the utility function is unbounded and has vanishing first derivatives (i.e., $\lim_{c \rightarrow \infty} u_i(c) = \infty$ and $\lim_{c \rightarrow \infty} u'_i(c) = 0$).

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 δ_i , or *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 Prize Lists. 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 prize 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. Evoking expected utility theory, the point $r_{i,t}^*$ below which the individual prefers early rewards and above which the individual prefers late rewards marks the point of indifference. Since no other periods are affected, indifference arises when

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

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

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

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 $r_\tau/r_{i,t}^*$ 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 convex, it is easy to highlight the bias alluded to in the introduction: consider two individuals with identical discount factors and preferences. Both obtain the same unemployment benefits b in the early period. But assume that person A has higher expectations about finding a job in the later period

than person B and therefore her income expectations are higher in the sense of first order stochastic dominance. Because the utility function is concave, $Eu'_i(y_{i,\tau})/Eu'_i(y_{i,t})$ is lower for person A than for person B , but then by (2) the measured ratio $r_\tau/r_{i,t}^*$ is higher for person A . As simple numerical example, consider an identical square root utility function, identical discount factor $\gamma_{i,\tau,t} = 1/2$, and $y_{A,1} = y_{B,1} = b = 1$, $y_{A,2} = 4$, $y_{B,2} = b = 1$, $r_\tau = 3$ and no uncertainty. Person A believes she will find a high-paying job for sure, while person B believes she will remain unemployed for sure. Using the equality in (2) then reveals that $r_{A,2}^* = 12$ while $r_{B,2}^* = 8$.⁵ Here person A looks more impatient in the sense of requiring a higher reward to wait to the future, and she is also the person who finds a job more quickly.⁶ But this does not indicate that more impatient people find jobs more quickly. In fact, she is *equally* impatient, and the mere fact that she knows that she will have a job *makes* her less interested in money in the future. 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. Consider the following alternative elicitation method where the individual returns are probabilistic, i.e., where

⁵We obtain $r_{A,2}^* = 12$ while $r_{B,2}^* = 8$ by solving the explicit version of (2):

$$\frac{1}{2} = \gamma_{i,\tau,t} = \frac{u_i(b + r_\tau) - u_i(b)}{Eu_i(y_{i,t} + r_{i,t}^*) - Eu_i(y_{i,t})} = \frac{\sqrt{4} - \sqrt{1}}{\sqrt{y_{i,t} + r_{i,t}^*} - \sqrt{y_{i,t}}}.$$

⁶The Multiple Price List method would provide an estimate for the discount factor of individual B of 3/8 (fairly close to the true value 1/2), while as a result of fluctuating income the estimate for individual A would be 1/4.

they receive a large amount R of additional income with some small probability, akin to a lottery win. They can choose to get some lottery tickets in the early period, or a larger number of lottery tickets in the later period. In the implementation we choose lottery tickets that can be "scratched" 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, the individual can have this additional income with probability ε_τ in the earlier period τ and with probability ε_t in the later period t . 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

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

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

$$\begin{aligned} \gamma_{i,\tau,t} &= \frac{\varepsilon_\tau}{\varepsilon_{i,t}^*} \frac{Eu(y_{i,\tau} + R) - Eu(y_{i,\tau})}{Eu(y_{i,t} + R) - Eu(y_{i,t})} \\ &\approx \underbrace{\varepsilon_\tau / \varepsilon_{i,t}^*}_{\text{probability ratio}} \quad \text{for large } R \text{ or } Eu(y_{i,t}) \approx Eu(y_{i,\tau}). \end{aligned} \quad (3)$$

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 $(u(R) - Eu(y_{i,\tau})) / (u(R) - Eu(y_{i,t}))$ because of vanishing marginal utilities, and this is bounded between $(u(R) - u(\bar{y})) / (u(R) - u(\underline{y}))$ and $(u(R) - u(\underline{y})) / (u(R) - u(\bar{y}))$, 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 of the late winning probability when the reward becomes sufficiently large, independent of beliefs about other income.*

3 Extensions

Our simple analysis relies on linearity in probabilities, and one might wonder how it would change if individuals weighted probabilities non-linearly in their utility. The simple analysis also assumed that an individual who wins the lottery early cannot save any of its proceeds to later periods, which seems restrictive after large lottery wins. Finally, one might ask how the model could be extended to encompass job search and job search effort, or more general savings choices. These are discussed in turn.

3.1 Non-expected utility

Our analysis relies on expected utility theory. A particular 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.

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.3 we show that we now obtain instead of (3) the adjusted formula

$$\gamma_{i,\tau,t} \approx \frac{w(\varepsilon_\tau)}{w(\varepsilon_{i,t}^*)} \approx 1 - \phi + \phi \frac{\varepsilon_\tau}{\varepsilon_{i,t}^*}, \quad (4)$$

where the first approximation holds for R large and the second for $\varepsilon_\tau/\varepsilon_{i,t}^*$ close to unity. Therefore, the discount factor is given approximately by an affine transformation of the probability ratio. While the probability ratio now does not directly uncover the discount factor, clearly individuals with higher revealed probability ratio have a higher discount factor. Therefore, the 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. Since conditional on winning the lottery we are in the gain domain, it can easily be shown that such reference dependence does not further affect our results.

3.2 Savings after lottery win

Our baseline model that equates the discount factor with the probability ratio in (3) relies on hand-to-mouth consumers. While this might be a palatable assumption for regular incomes, it might be especially contentious in cases of large lottery wins. Consider therefore the same setting as before but abstract from uncertainty (though income differences by period and person remain) to create an environment where the savings game among sophisticates is well understood. Now add the simplest of saving choices where individuals can save at interest rate ι after a lottery win as long as their net worth remains positive, i.e., wealth and net present value of future earnings remains above zero. That means that winning early allows them to benefit from this for several periods.

For this extension, we focus only on exponential discounters ($\gamma_{i,\tau,t} = \delta_i^{t-\tau}$) or hyperbolic discounters (identical, except that $\gamma_{i,0,t} = \beta_i \delta_i^t$), and 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 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 hyperbolic discounters, and they play a game with their future selves. In the latter case we restrict attention to Bernoulli utility with constant relative risk aversion above unity, for which this savings game is well understood (Laibson (1996)). The equilibrium for finitely-lived individuals is unique, and we focus on the limit of these when we study infinitely lived individuals. 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 $c_{i,t} = y_{i,t}$.

Denote by $U_{i,\tau}(R + NPV_{i,\tau})$ 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 $NPV_{i,\tau}$ 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

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

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 $U_{i,\tau}(R + NPV_{i,\tau})$, 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 probability $\varepsilon_{i,t}^*$ where both terms are equalized, which after rearrangement yields question:⁷

$$\begin{aligned} \gamma_{i,\tau,t} &= \frac{\varepsilon_\tau U_{i,\tau}(R + NPV_{i,\tau}) - V_{i,\tau}}{\varepsilon_{i,t}^* U_{i,t}(R + NPV_{i,t}) - V_{i,t}} \\ &\approx \frac{\varepsilon_\tau U_{i,\tau}(R + NPV_{i,\tau})}{\varepsilon_{i,t}^* U_{i,t}(R + NPV_{i,t})} \text{ for } R \text{ large} \\ &\approx \frac{\varepsilon_\tau U_{i,\tau}(R)}{\varepsilon_{i,t}^* U_{i,t}(R)} \text{ for } R \text{ large.} \end{aligned} \tag{5}$$

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;

⁷The indifference condition is

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

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

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 $NPV_{i,\tau}$ and $NPV_{i,t}$ 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.⁸ It is also well-known for standard cake-eating problems with infinitely-lived exponential discounters⁹, which carry over to naive hyperbolics who expect to be exponential after the initial period. For infinitely-lived sophisticated 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 consumers:

Proposition 2 *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*

$$\gamma_{i,\tau,t} \approx \frac{\varepsilon_\tau}{\varepsilon_{i,t}^*} \quad (6)$$

as in (3) for sufficiently large lottery rewards if (i) the individual is an ex-

⁸While this might be less obvious in the game between selves of a sophisticated 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.

⁹For 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 hyperbolics from the second period onward, a similar logic applies to them.

ponential discounter, (ii) the individual is a sophisticated or naive hyperbolic discounter and $\tau > 0$.

In the case of hyperbolic discounters and immediate rewards ($\tau = 0$) the probability ratio understates 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 (5) converts to (6) 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 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 (6) 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.

3.3 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 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. 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 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 2 we obtain:

Proposition 3 *Consider infinitely-lived naive 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 envelop theorems: consider the consumption sequence $C_{i,1}, C_{i,2}, C_{i,3}, \dots$ of individual i in the ab-

sence 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 hyperbolic individuals or exponential discounters are characterized by a standard inter-temporal optimization problem, and we can rely on the envelop 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 $C_{i,1}, C_{i,2}, C_{i,3}, \dots$ 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 $C_{i,1}, C_{i,2}, C_{i,3}, \dots$, 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 envelop theorem does not apply, as their future consumption is determined by a game against 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.

3.4 A general model with job search and savings.

Finally, our analysis can be embedded into a richer overall model. In the Additional Material we outline how job search effort, savings decisions throughout the life, and job acceptance rates can be embedded into this setting in an abstract way. This a demonstration of how such a setting can capture for example the job search model by DellaVigna and Paserman (2005). Similar to the previous argument, at least for exponential discounters and naive hyperbolics the envelop theorem allows us to arrive at (5).

4 Validation Experiment

4.1 Basic setup

We validate our method experimentally on two student samples of around 100 students each. The experiment was conducted on-line and was pre-registered on the AEA RCT Registry (Belot et al. (2019a)). The subject pool consists of first and second year students at the School of Business and Economics of the VU University Amsterdam. Participants were recruited through the 'Research Participation System' and participated through a Qualtrics survey.

In order to contrast a "stable" situation where no shocks are expected to a situation where changes in income or expenditures are expected, we recruited two waves at two different points in time: The first wave was recruited in November 2019, before the two major December holidays in the Netherlands (Saint Nicholas and Christmas), that is, at a time when their current budget is reduced by extraordinary expenditures for gifts. The second wave was recruited in February 2020 when no such extra constraints exist (stable scenario). Note that the threat of Covid 19 was not yet present in Europe at that time.

Participants were recruited in the exact same way and drawn from the same student pool. There was no mention of a second wave when recruiting participants for the first wave. Those participating in the first wave were automatically excluded from participation in the second wave.

4.2 Experimental tasks

Participants were asked to complete 4 consecutive blocks of questions: (i) a block of 20 convex-budget-set (CBS, Andreoni and Sprenger (2012)) questions, (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 order of (i) and (ii) was randomized.

The CBS and MLL questions have equivalent rates of return ranging from 20% to 100%. Early options are either the week of the experiment or 8 weeks after it, and late options are either 5 weeks later or 14 weeks later. CBS questions consist of an early and a late option that vary in time of payout and token-exchange-rate; each participant has 100 token per question to allocate between the two options. MLL questions consist of an early and a late option that vary in the number of lottery-scratch-cards and timing; each participant chooses for each question either the early or late option. One question out of all CBS and MLL questions is randomly selected for payout. For instructions and complete questions see the Supplementary Material.

The selected parameters for the blocks of questions for CBS and MLL are summarized in Table 1. Each row is one question, and specifies t (the early period), k (delay, difference between the early and late period in days), a_t (token exchange rate early period for CBS, number of early period lottery tickets for MLL), a_{t+k} (same for late period) and $1 + r$ (implied exchange rate). We randomized at the individual level whether *all* question blocks had ascending interest rates or *all* blocks had descending interest rates.

At the end of the questionnaire participants provided their contact details and bank accounts. They received an email on the same day stating the randomly selected question and amount(s) of money they would receive on particular days. Each participant received 4 Euro for participation plus an amount based on the selected question (money or lottery tickets). The participation fee plus any money resulting from the ‘Today’ choices was transferred and received on the same day (transfers between banks within the Netherlands occur immediately).

Table 1: Question parameters

Convex budget set (CBS) questions					Lottery ticket (MLL) questions				
t	k	a_t	a_{t+k}	$1+r$	t	k	a_t	a_{t+k}	$1+r$
early period	delay	early rate	late rate		early period	delay	early # of tickets	late # of tickets	
0	35	0.1	0.12	1.2	0	35	5	6	1.2
0	35	0.1	0.14	1.4	0	35	5	7	1.4
0	35	0.1	0.16	1.6	0	35	5	8	1.6
0	35	0.1	0.18	1.8	0	35	5	9	1.8
0	35	0.1	0.2	2	0	35	5	10	2
0	98	0.1	0.12	1.2	0	98	5	6	1.2
0	98	0.1	0.14	1.4	0	98	5	7	1.4
0	98	0.1	0.16	1.6	0	98	5	8	1.6
0	98	0.1	0.18	1.8	0	98	5	9	1.8
0	98	0.1	0.2	2	0	98	5	10	2
56	35	0.1	0.12	1.2	56	35	5	6	1.2
56	35	0.1	0.14	1.4	56	35	5	7	1.4
56	35	0.1	0.16	1.6	56	35	5	8	1.6
56	35	0.1	0.18	1.8	56	35	5	9	1.8
56	35	0.1	0.2	2	56	35	5	10	2
56	98	0.1	0.12	1.2	56	98	5	6	1.2
56	98	0.1	0.14	1.4	56	98	5	7	1.4
56	98	0.1	0.16	1.6	56	98	5	8	1.6
56	98	0.1	0.18	1.8	56	98	5	9	1.8
56	98	0.1	0.2	2	56	98	5	10	2

Those that earned lottery tickets received scratch tickets by postal mail on the corresponding date. Each ticket had a 1 Euro buying price and could be scratched immediately to discover whether it paid a price. The tickets pay out a "headline" price of 25.000 Euro with a chance of approximately 1 in 2 million, but they also offer smaller prices (at larger odds).

Thus, within method the transaction costs are kept constant across choices at different time horizons: monetary payoffs are transmitted by bank transfer, lottery tickets are provided through mail delivery. That means that individuals who choose to receive payoff in the week of the experiment receive them with slight delay (bank transfers and mail occur on the same day, but mail is delivered the morning after).

4.3 Descriptive statistics

Table 2 presents basic descriptive statistics for each sample. Despite being recruited in the same way, we observe some differences in observable characteristics. Not surprisingly, students in wave 2 are slightly older than students recruited in wave 1 ($p=0.07$). But we also have a higher proportion of foreign nationals in the first wave (14% in Wave 1, 5% in Wave 2, $p=0.02$) and significant differences in shares of men (57% in Wave 1 and 80% in Wave 2, $p=0.00$). We do not have a good explanation for why more male students registered in the second wave.¹⁰ Despite these differences, the average monthly income is similar at 740 Euro in Wave 1 and 732 Euro in Wave 2. We will address below concerns about how differences across these two samples affect our analysis.

We also collected information about expected future income and expenditure changes, distinguishing between gift and non-gift related changes for participants in Wave 1. Expectations about non-gift related income changes in the next three months are similar between waves (around 70% expect no changes, 24% expect an increase and the rest expects a decrease). Regarding expectations about non-gift related expenditure changes, the fraction that expects no change is roughly 60% in both samples, but for the rest the distribution is significantly different: Participants in wave 1 are less likely to expect a decrease and more likely to expect an increase. Such difference is not necessarily unexpected; for example expenditures related to travel or holidays may also increase for participants in wave 1. Participants in Wave 1 expect to spend an *average* of 90 Euro on Christmas or Santa Claus while subsequently they expect to receive on average 65 Euro in monetary gifts (see Table 3).

Using the information on expectations, we classify individuals into those expecting a ‘shock’ and those not expecting a shock. The ‘no shock’ group is defined as individuals that indicate (1) no expected changes in expenditure (2) no expected change in income (3) minor expenditures for *and* income from Santa-Claus/Christmas presents (below 50 Euro in such expected expenditures

¹⁰There are very few studies noting gender differences in measured time preferences. We only know of one study documenting that women were found to be more patient than men based on a multiple price list experimental measure (Dittrich and Leipold (2014)).

Table 2: Descriptive Statistics

	(1)		
	Wave 1	Wave 2	P-val
Demographics:			
Age	19.0	19.4	0.07
Male (%)	0.57	0.80	0.00
Country of origin			
Netherlands (%)	0.86	0.95	0.02
European country (%)	0.07	0.03	0.13
Other (%)	0.07	0.02	0.10
Income/Expenditure shocks			
Income per month	740	732	0.91
Expected income decrease (%)	0.04	0.05	0.72
Expected no income change (%)	0.74	0.71	0.73
Expected income increase (%)	0.23	0.24	0.84
Expected expenditure decrease (%)	0.03	0.10	0.04
Expected no expenditure change (%)	0.59	0.63	0.52
Expected expenditure increase (%)	0.39	0.28	0.09
Any expected shock (income/exp.) (%)	0.84	0.52	0.00
Observations	108	105	

Table 3: Holiday gift expenditures/receipts

	(1)
	Wave 1
	mean
Spending on Santa Claus presents	24.2
Spending on Christmas presents	66.7
Receiving as Santa Claus presents	11.6
Receiving as Christmas presents	54
Net holiday expenditures	25.3
Observations	106

and in such expected gifts). The ‘shock’ sample contains all other participants. Clearly for Wave 2 only the first and second criterion apply and 52% are labeled as expecting a shock, while the additional expenditures expectations for Christmas increase this to 84% in Wave 1 (Table 2).

The average responses to the CBS and MLL questions are presented in Figure 2 in Online Appendix B.1. Mean responses are in line with expectations: patient choices increase with the interest rate and decrease with the delay length. On average, the two waves yield similar pictures, although there seems to be slightly more present bias in Wave 1.

4.4 Comparing CBS and MLL

We proceed by estimating preference parameters for each individual using responses to both the CBS and MLL questions. Our MLL questions allow bounds identification of δ and β , equivalent to the conventional Multiple Price List elicitation method. As outlined in our pre-analysis plan, we use questions with 5 weeks between early and late options to identify our MLL parameters, as this serves as a validation to our experiment with job seekers presented in the next section that used similar questions.¹¹ MLL choices then identify whether the discount factor falls into one of six intervals. As is common in the literature, we pick the midpoint of the identified interval for each parameter estimate.¹² Andreoni and Sprenger (2012) show that point estimates are identified from the CBS responses if an individual made at least two interior choices (as opposed to corner solutions), using a two-sided Tobit regression.¹³ The distributions of estimated preference parameters are shown in Online Appendix B.1 in Figure 3 (Wave 1) and Figure 4 (Wave 2), separately by method.

¹¹We included questions with 14 weeks delay because the CBS method relies on a larger set of questions to identify individual discount factors. Obviously answers to these questions also provide additional information for MLL. The Online Appendix C.2 shows very similar insights when we use all questions to bound individual discount factors.

¹²Both estimation procedures are presented in more detail in the Online Appendix C.1.

¹³This works for half of the sample. For one-third we use simple OLS (as Andreoni and Sprenger (2012) do in their appendix Table A4). The remaining participants put 100% of their tokens in the patient choice, *for all* questions. As a result only a lower bound for their time preferences δ and β is identified, and for these individuals we set $\hat{\delta} = \hat{\beta} = 1$.

CBS measures are more continuous than the MLL measures. For comparability we also construct “coarse” measure of CBS, where we transform the values into the six MLL values, based on the interval in which they fall. The pre-analysis plan specified the following hypothesis for the comparison between MLL and CBS:

Hypothesis 1: In Wave 2, where income and expenditure shocks are rare, we expect a similar ranking of individuals: individuals with a higher measurement under MLL should also have also a higher measurement under CBS. For Wave 1 we expect a weaker association between the two.

While this is our main hypothesis, we also specified further analysis along the following lines as sub-ordinate exercises:¹⁴

Sub-Hypothesis 1.1: In Wave 2, CBS and MLL provide similar magnitudes for the discount factor: the differences in discount factors between individuals should be similar, and possibly the level of the discount factor. In wave 1 again we do not expect this.

Sub-Hypothesis 1.2: For individuals who report no shocks in expenditures and/or income we expect a higher discount factor under MLL to be associated with a higher discount factor under CBS (in either wave). We do not expect the same for those with shocks.

Hypothesis 1 is based on the robust theoretical result that higher discount factors are linked to higher MLL measures. In wave 2 where shocks are low the CBS method should uncover discount factors, and so higher CBS should be indicative of higher MLL. Note that we expect a weaker association in Wave 1 for both δ and β , as fluctuations surrounding Christmas are likely to affect even answers to the future questions (e.g., 8 weeks versus 13 weeks), even though

¹⁴In our pre-analysis plan Sub-Hypothesis 1.1 was called "Hypothesis 1a". Sub-hypothesis 1.1 was directly mentioned as an extension analysis to the main Hypothesis 1.

both of these weeks are *after* Christmas. For example, large expenditures in December may make income in January more valuable for some. Alternatively, some individuals may expect large monetary gifts during Christmas, lowering the value of income in January. Since methods differ in how they deal with shocks (CBS is not designed to be robust against such shocks), we expect them to correlate less well in these settings. Sub-hypothesis 1.1 states that magnitudes should be comparable, which is only true in our baseline theory or in the case of savings for the long-run discount factor. Finally, Sub-hypothesis 1.2 restates the same on an individual-level basis, where CBS should be a good comparison for those without shocks but not otherwise.

To test these, we regress our MLL estimates on the CBS estimates:

$$\hat{\delta}_i^{MLL} = \gamma_0 + \gamma_1 \hat{\delta}_i^{CBS} + \varepsilon_i \quad (7)$$

and similar for β . The formal tests in our pre-analysis plan are then:

- Hypothesis 1: $\gamma_1 > 0$ in Wave 2, but not necessarily in Wave 1.
- Sub-hypothesis 1.1: $\gamma_1 = 1$ in Wave 2, but not necessarily in Wave 1, using coarse CBS, especially for δ . Alternatively: test equality of distribution of δ and β in Wave 2, using coarse CBS.
- Sub-hypothesis 1.2: $\gamma_1 > 0$ for those without shocks in either wave.

Importantly, testing these hypotheses does not require that both samples are similar in observable and unobservable characteristics. Thus, the differences that may exist between the two samples do not matter for this analysis.

Regression results by wave are presented in Table 4. The estimates for δ (columns 1 and 2) show that the association is positive in both waves, but indeed only statistically significant in Wave 2. The coefficient is also larger in Wave 2 (0.588) than in Wave 1 (0.244). For our estimates of β (columns 3 and 4), the association between the two methods is weaker in general, but again largest and statistically significant in Wave 2. These findings support Hypothesis 1: MLL picks up a similar variation in the discount factor as CBS in the sample with a higher fraction of individuals without income/expenditure

shocks, i.e., a higher fraction of individuals that fulfill the assumptions underlying CBS.

Table 4: Regression MLL on CBS

	Wave 1 (1) Delta (MLL)	Wave 2 (2) Delta (MLL)	Wave 1 (3) Beta (MLL)	Wave 2 (4) Beta (MLL)
Delta (CBS)	0.244 (0.148)	0.588*** (0.208)		
Beta (CBS)			0.110 (0.077)	0.220*** (0.077)
Constant	0.594*** (0.136)	0.221 (0.193)	0.859*** (0.078)	0.785*** (0.079)
N	108	105	108	105

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficients are not equal to unity, but this might simply be due to the coarseness of the MLL measure. For that reason our pre-analysis plan specified to test Sub-hypothesis 1.1 with coarse CBS.¹⁵ If we compare MLL to this coarse CBS measure using the same regression, we obtain even stronger results (Table 8 in the appendix): the coefficient for δ in Wave 2 is 0.901, highly statistically significant and not statistically different from 1. In Wave 1 it is only 0.383 (significant at the 10% level). Also for β we find a significant coefficient in Wave 2 (0.388) and a smaller non-significant coefficient in Wave 1 (0.191). This supports the first test for Sub-hypothesis 1.1, at least for the long-run discount factor, indicating that even the magnitude of differences across individuals is meaningful in MLL. It does so even though the two methods deliver different means of δ and β . Indeed, a paired t-test for equality of the two measures (using coarse CBS) rejects this hypothesis: both average δ (p-value < 0.01) and average β (p-value = 0.04) differ significantly in Wave 2.

¹⁵An alternative is to consider the rank of the estimates within the Wave (per method). In Online Appendix C.3 we reproduce our main analysis using ranks rather than levels and find that results are very similar. This was not part of our pre-analysis plan.

A comparison of the two waves confirms our main hypothesis that MLL correlates well with CBS when a small fraction of individuals is anticipating shocks, but does not correlate well when a larger fraction of individuals are expecting shocks, that is, in settings in which the basis for CBS (and most other methods) is violated.

So far the analysis was run on each wave separately, including all individuals irrespective of whether they reported expecting a change in income/expenditures or not, exploiting the variation due to the timing of the two waves.

But we do not need a source of exogenous variation in shocks to test these hypotheses. Whether the shock is exogenous or not, i.e. whether it is correlated with unobservables or not, we would again expect a stronger correlation between MLL and CBS for those who do not expect shocks than for those who do.

To isolate more clearly the role of shocks as the factor causing the difference, we now turn to *individual information* on income and expenditure shocks to test Sub-hypothesis 1.2. We classify individuals from both waves as those with or without shocks (as described above) and estimate equation (7) for the two groups. Results for δ are presented in Table 5. We find again that the measures correlate much more strongly in the samples without shocks: for both waves we obtain a highly significant coefficient for the subgroups without shocks (columns 1 and 3). Both coefficients are also very close to 1, that is, the methods do not only capture the correct ranking but also the correct magnitudes of differences between individuals (see Sub-hypothesis 1.1). On the contrary, we find insignificant and small coefficients for the samples that report expected shocks in column 2 and 4.

A similar breakdown for β is presented in Table 6. As the association in general is much weaker for β , the point estimates provide little information regarding differences in alignment for individuals with and without shocks. The coefficients have equal magnitude, but only for the no-shock sample it differs significantly from zero (column 3).

We also provide tables similar to Table 5 and 6 using coarse CBS in On-

line Appendix B.1 (see Tables 9 and 10). Results are overall rather similar, except that the coefficient on β for no-shock individuals in wave 2 is much larger (0.696), significantly different from zero, and insignificantly different from unity.

Table 5: Regression MLL estimates on CBS estimates by income/expenditure shocks: Delta

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Delta (CBS)	1.248*** (0.342)	0.092 (0.164)	1.048*** (0.309)	0.284 (0.283)
Constant	-0.307 (0.302)	0.737*** (0.152)	-0.211 (0.289)	0.501* (0.259)
N	17	89	50	55

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Regression MLL estimates on CBS estimates by income/expenditure shocks: Beta

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Beta (CBS)	0.132 (0.107)	0.102 (0.090)	0.201** (0.092)	0.214 (0.191)
Constant	0.836*** (0.112)	0.865*** (0.090)	0.843*** (0.101)	0.755*** (0.185)
N	17	89	50	55

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

These results confirm our main hypothesis (Hypothesis 1) about the comparison between MLL and CBS: both measures correlate strongly when few people expect shocks (stable scenario) but not otherwise. This is related to

the shocks individuals perceive, as within wave the strong correlation exists for those without shocks (Sub-hypothesis 1.2). Both methods uncover similar differences in long-run discount factors between individuals (first test of Sub-hypothesis 1.1) even though they do not uncover exactly the same distributions (second test) in the stable scenario. These results indicate that our main premise is justified: MLL is suitable to measure which individuals are more patient, based on a similar performance to CBS in low-shock samples. But it is designed for scenarios with shocks that are not accounted for by CBS and most other methods. In such scenarios, we find a different ranking of individuals using the MLL method.

Our pre-analysis plan also specified similar tests with a group of job seekers for the late spring 2020, but due to Covid restrictions the recruitment for these could not be carried out. It also included a second main hypothesis intended to investigate each method across samples under the assumption that “the underlying distribution of time preferences of individuals in sample 1 and sample 2 are similar” (Belot et al. (2019a)). If the only difference is that those who answer the questions in Wave 1 have additional expenditures currently but more income in the future compared to individuals in Wave 2, we expected CBS to show more present bias in Wave 1 than in Wave 2, and not MLL. MLL should be shielded from income/expenditure changes, while CBS should pick up the current shortness of money and the future abundance in Wave 1. But this is unfortunately not the only difference between waves. As discussed earlier, observables differ significantly across waves, casting doubt on whether the unobservables, such as true discount factors, are distributed identically across waves. Since this exercise depends on similar distributions, we do not present this analysis here and relegate it to Online Appendix B.2. We find more present-bias in Wave 1 but estimates of the difference are not

significantly different between methods.^{16,17}

5 Application to Job Search

We apply our new 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)). We are not aware of any study that takes direct measures of time preferences and links them to job search activities of the unemployed, maybe because of the difficulties of varying incomes discussed in the introduction. We fill this gap by illustrating the application of our method in this setting.


We collected information on discount factors of individuals 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.¹⁸ 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.

¹⁶Participants in the two waves were drawn from the same pool of students and both studies filled up quickly, suggesting abundant supply of students. Nevertheless, there are factors that may explain differences in observables and discount factors. First, all participants from Wave 1 were automatically excluded from participation in Wave 2. The most ‘eager’ individuals may thus have ended up in the first sample, while the remaining less eager individuals are in Wave 2. Second, while all students in the pool were required to participate in some experimental studies, they were free to decide when to do so. Students signing up in the first semester may be different types of students from those signing up in the second semester.

¹⁷Note that neither elicitation method correlates well with observables or ex-post collected personality traits, and the few correlations that are detected in one wave are generally not stable across waves.

¹⁸Lotteries again paid out according to the first ten rows of Table 1, except that the delay was four weeks (28 days) instead of five weeks (35 days). We changed this for the validation experiment by one week to avoid payouts during the holiday period between Christmas and New Year in Wave 1. Lottery tickets were of comparable value, costing £2 in the UK. They had a headline prize of £250000 and again also lower prizes.

Figure 1: Estimated time preference parameters job seekers



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. The 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.

5.1 Time preference parameters

We start by presenting descriptive statistics on the elicited time preference parameters in the standard hyperbolic discounting (beta-delta) framework. In line with the analysis above, for each individual we infer present-bias (β) and long run discount factors (δ). Figure 1 shows the distribution of δ (calculated

Table 7: Regressions job search effort and log reservation wage

	(1)	(2)	(3)	(4)	(5)
	Weekly hours spent on job search	Weekly applications	Min. (hourly min. wage)	Log hourly stated res. wage	Weekly interviews
Beta	0.731 (1.407)	0.119 (0.163)	0.006 (0.009)	0.054 (0.081)	0.645*** (0.287)
Delta	-3.731* (2.250)	0.025 (0.265)	0.001 (0.020)	-0.007 (0.110)	-0.668*** (0.116)
Model	Linear RE	Poisson RE	Linear RE	Linear	Poisson RE
N	383	530	516	166	490

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.

based on choices of payments in 8 vs 12 weeks) and β . We find that individuals display a bimodal distribution of long-run discount factors. The bimodal distribution arises because a large fraction of individuals always choose the early (or late) date. Regarding β , we find that 71% are time-consistent ($\beta = 1$) and 20% are present biased.

5.2 Time preferences, search effort, and job interviews

We have two 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.¹⁹

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 low. 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 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.

In Table 7 we show that less present bias indeed significantly increases job interviews (column 5), which we take as a close proxy for their re-employment chances.²⁰ All regressions also contain standard controls (not reported), including education, age and gender. Point estimates for hours spent searching (column 1) and number of applications (column 2) are positive but small in magnitude and not statistically significant. For the reservation wage, the results are shown in column 3 and 4. Effects of present bias on either measure of the reservation wage are economically small and statistically insignificant.

¹⁹Note 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.

²⁰Presented regression results are from a Poisson regression (since interviews is a count variable). We observe the outcome 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.

In sum, these findings seem broadly in line with the predictions of 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 7 reveals that those with a higher long-run discount factor have significantly fewer job interviews. They also have significantly lower hours of job search, while we find no relationship with job applications. There is no economic or statistically relevant relationships with either measure of reservation wages. These findings are not as predicted by DellaVigna and Paserman (2005).

5.3 Discussion of job search and time preferences

The findings on present-bias are broadly encouraging for job search models that include impatience as a driving force. The findings on the long-run discount factor are also 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 are more choosy on 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. 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.

6 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 with students, the discount factors elicited with this method correlates well with those from standard (convex budget set) elicitation methods in the absence of shocks. Our measure is intended to absorb shocks to income while others are not, and indeed shocks break this strong correlation.

An illustration with job seekers showcases how our method might be applied in settings where income shocks are the norm, and vindicates previous predictions on present-bias while raising new aspects related to the long-run discount factor.

Future work might explore more deeply how our method performs relative to, say, structural estimation techniques that incorporate elicitation of subjective expectations about future variation in income and consumption. This is an open agenda. One main advantage might be that our method is simple

to administer and analyze, and does not require strong modelling assumptions regarding utility functions or the space of beliefs.

References

- Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2008). Eliciting risk and time preferences. *Econometrica*, 76(3):583–618.
- 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., Niederle, M., and Sprenger, C. (2015). Working over Time: Dynamic Inconsistency in Real Effort Tasks *. *The Quarterly Journal of Economics*, 130(3):1067–1115.
- Belot, M., Kircher, P., and Muller, P. (2019a). Eliciting time preferences using high stakes lottery tickets. *AEA RCT Registry, No AEARCTR-0005115*, <https://doi.org/10.1257/rct.5115-1.0>.
- 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.
- 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.
- Dittrich, M. and Leipold, K. (2014). Gender differences in time preferences. *Economics Letters*, 122(3):413 – 415.
- Dynarski, S. and Gruber, J. (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity*, 28(1):229–303.
- Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424.
- 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.
- 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.
- 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.
- Meyer, C. J. (2018). In search of a better life: Self-control in the ethiopian labor market.
- Morand, O. F., Reffett, K. L., and Tarafdar, S. (2015). A nonsmooth approach to envelope theorems. Ssrn working paper no 2609292.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343.
- 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.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.

A Online Appendix: Theory - Main Model Extensions and Illustrations

A.1 Illustration with real effort elicitation

Assume individuals value money and leisure. In the simplest setting their utility function per period can be represented as $E[u_i(c_{i,t}) + v_i(h_{i,t})]$, where $u_i(\cdot)$ and $y_{i,t}$ are as in the main body, and $v_i(\cdot)$ is an increasing and convex utility-of-leisure function and $h_{i,t}$ is the leisure in the given period. If individuals can choose between r_τ units of time spent on a real effort task in period τ , and r_t units of real effort to spent on a task in period $t > \tau$ (and they might get some additional payment in the first period as a show-up fee to make them willing to take either one of these), their level of indifference is now given by

$$Ev_i(h_{i,\tau} - r_\tau) + \gamma_{i,\tau,t}Ev_i(y_t) = Ev_i(h_{i,\tau}) + \gamma_{i,\tau,t}Ev_i(h_{i,t} - r_{i,t}^*). \quad (1')$$

In this simple example we obtain the analogous equation to (1), only that we are now working with leisure time rather than income. So we can again derive the analogue to (2)

$$\gamma_{i,\tau,t} = \frac{Ev_i(h_{i,\tau} - r_\tau) - Ev_i(y_{i,\tau})}{E v_i(h_{i,t} - r_{i,t}^*) - Ev_i(y_{i,t})} \approx \frac{r_\tau}{r_{i,t}^*} \frac{Ev'_i(h_{i,\tau})}{Ev'_i(h_{i,t})}, \quad (2')$$

where the third expression uses the Taylor approximation $v_i(h_{i,t} - r) \approx v_i(y) - v'_i(y)r$.

Now a similar example as that in the main body goes as follows: Consider two individuals A and B with identical utility-of-leisure functions $v_i(h) = \log(h)$ and identical discount factor, who differ only in their expectations about finding a job and associated changes in leisure time: They are both unemployed in the early period with full leisure time $h_{A,1} = h_{B,1} = 1$. In the late period individual A believes to have found a job that cuts her leisure to $h_{A,2} = 1/4$, while individual B believes that she will stay unemployed with $h_{B,2} = 1$. Assume real effort r_1 is sufficiently small that the approximation in (MPL') is

valid and we get

$$\begin{aligned} \gamma_{A,1,2} &= \frac{r_\tau}{r_{A,t}^*} \frac{Ev'_i(1)}{Ev'_i(1/4)} = \frac{r_\tau}{4r_{A,t}^*} \\ &= \gamma_{B,1,2} = \frac{r_\tau}{r_{B,t}^*} \frac{Ev'_i(1)}{Ev'_i(1)} = \frac{r_\tau}{r_{B,t}^*}, \end{aligned}$$

where the two lines equal because of the assumption of identical discount factor across individuals. So $\frac{r_\tau}{r_{A,t}^*} = 4\frac{r_\tau}{r_{B,t}^*}$, and if one neglects the change in expected leisure time person A appears as if she has a higher discount factor than person B . Person A also finds a job more quickly. If one just took $r_\tau/r_{i,t}^*$ as a measure of patience, the more patient person here finds the job more quickly. But truly they both hold the same discount factor and person A only does effort early because she anticipate a reduced leisure endowment in the second period.

A.2 Illustration of direct method elicitation

The direct method of Attema et al. (2016) asks questions of the kind: Do you want to get 10 Euro in each of the first X weeks, or would you rather get 10 Euro in each week after X until some maximum week. Let the maximum week for simplicity be week ten. Clearly if $X = 9$ most people would like to be paid early since they would be paid nine periods while they will only be paid in one period if they take the late payments. Similarly, if $X = 1$ most people would presumably take the late payments as it will lead to nine payments and despite discounting this will for most be preferred to a single early payment. Varying X elicits the point where the individual is indifferent between early and late payment. The discount factor can be recovered from versions of such questions when background consumption does not change. People with identical discount factors have to answer the questions identically.

But now consider two people that have identical discount factor and preferences but different job finding expectation. Person A expects to find a job in week 6. Person B does not expect to find a job within the ten weeks. Both get low unemployment benefits during unemployment and consume them. Upon

finding a job person A gets a higher wage w .

Consider first individual B: If she does not discount at all, she would choose indifference point $X=5$, as she would like to have the maximum number of periods payout. If she has mild time preferences, she might choose $X=4$ because early payments are valued more. Now consider individual A: Assume w is so high that the marginal utility $u'(w)$ is close to zero. So she essentially does not value additional money in periods after period 5. That means he would choose $X=2.5$ if she does not otherwise discount, or with some mild discounting she would choose $X=2$. So even with this method, it would look as if A is much more impatient.

A.3 Probability Weighting

Consider an individual who obtains $y_{i,t}$ in period t , unless she wins the lottery in which case she obtains $y_{i,t} + R$. We abstract from uncertainty in $y_{i,t}$ for illustrative purposes. Such uncertainty would need to be taken into account with appropriate probability weights, which does not alter the final result but significantly increases notational complexity. Let w be a mapping from $[0,1]$ onto $[0,1]$ representing the probability weighting function, and we adopt rank-dependence. Indifference between the early and late lottery now requires

$$\begin{aligned} & [(1 - w(\varepsilon_\tau))u(y_{i,\tau}) + w(\varepsilon_\tau)u(y_{i,\tau} + R)] + \gamma_{i,\tau,t}u(y_{i,t}) \\ = & u(y_{i,\tau}) + \gamma_{i,\tau,t}[(1 - w(\varepsilon_{i,t}^*))u(y_{i,t}) + w(\varepsilon_{i,t}^*)u(y_{i,t} + R)], \end{aligned}$$

where winning the lottery is always the most attractive outcome and is weighted by the probability weight. For the less attractive outcome rank-dependence means that it is assessed with $1 - w(p)$. This reduces to

$$\begin{aligned} \gamma_{i,\tau,t} &= \frac{w(\varepsilon_\tau)}{w(\varepsilon_{i,t}^*)} \frac{u(y_{i,\tau} + R) - u(y_{i,\tau})}{u(y_{i,t} + R) - u(y_{i,t})}. \\ &\approx \frac{w(\varepsilon_\tau)}{w(\varepsilon_{i,t}^*)} \text{ for large } R. \end{aligned} \tag{8}$$

The steps to show the approximation are identical to those used in (3). We assume that $\varepsilon_\tau R = K$ for some fixed K , and take limits as R becomes large.

Let $r_i = \varepsilon_\tau / \varepsilon_{i,t}^* \in (0, 1)$ be the limit of individual's choices as R becomes large. We can then derive the following approximation around r_i of unity:

$$\begin{aligned} \ln \left(\frac{w(\varepsilon_\tau)}{w(\varepsilon_{i,t}^*)} \right) &= \ln \left(\frac{w(\varepsilon_{i,t}^* r_i)}{w(\varepsilon_{i,t}^*)} \right) = \ln w(\varepsilon_{i,t}^* r_i) - \ln w(\varepsilon_{i,t}^*) \\ &\approx \ln w(\varepsilon_{i,t}^*) + \frac{w'(\varepsilon_{i,t}^*)}{w(\varepsilon_{i,t}^*)} (\varepsilon_{i,t}^* r_i - \varepsilon_{i,t}^*) - \ln w(\varepsilon_{i,t}^*) \\ &\approx \phi(r_i - 1) \end{aligned}$$

where the first approximation is simply a Taylor expansion and the second reflects that $w'(\varepsilon_{i,t}^*)\varepsilon_{i,t}^*/w(\varepsilon_{i,t}^*)$ captures the elasticity of the weighting function around zero, which we denote by ϕ . Therefore (8) can now be written as

$$\begin{aligned} \gamma_{i,\tau,t} &\approx e^{\phi(r_i-1)} \\ &\approx 1 + \phi(r_i - 1) \end{aligned}$$

where the second approximation reflects that standard approximation of exponential functions around an argument of zero.

A.4 Setup with Savings after Lottery Win

Here we outline the setting where individuals can save after their lottery win. Recall that individuals at time $t_0 = 0$ at time t have continuation utility

$$E \sum_{s=t}^T \gamma_{i,t,s}(t_0) u_i(c_{i,t}) \tag{9}$$

where total life span T could be infinite. (Note that a decision-maker at time t_1 might have a different continuation value because she applies discount factors $\gamma_{i,t,z}(t_1)$.) Individuals are hand-to-mouth unless they win the lottery ($c_{i,t} = y_{i,t}$). But in case of a lottery win they can freely save or borrow at

interest rate ι as long as their wealth W (i.e., the net present value of past savings plus current and future income) stays weekly positive. So starting a period with W allows consumption

$$c \in [0, W] \tag{10}$$

and next period wealth

$$W' = (1 + \iota)(W - c). \tag{11}$$

In particular, this means that individuals can use proceeds from lottery wins over many periods. For ease of exposition also assume that the income stream is deterministic, though possibly heterogeneous across individuals and time.

Here we focus only *exponential* discounters or *hyperbolic* discounters as defined in the setup of the basic model. We assume $\delta_i(1 - \iota) \leq 1$ so exponential discounters dis-save after a lottery win. In the hyperbolic case, we follow the literature and distinguish between *naive* individuals who believe that they will behave as exponential discounters in the future, and *sophisticated* individuals who understand that in their future "selves" will also have more interest in immediate consumption. Both exponential and naive individuals believe that any sequence of savings choices that is optimal for them today will also be optimal for their future selves, so their savings problem after a lottery win is simply an optimization problem: choose sequence c_t, c_{t+1}, \dots to maximize (9) subject to constraints (10) and (11). Sophisticates on the other hand understand that future selves discount the future different from themselves and will take different actions from the ones that the individual would like to commit to today. It plays a game with its future selves, as, e.g., in Laibson (1996).

For finitely-lived individuals the savings problem has a unique solution. In the case of sophisticates it is found by backward induction: under constant relative risk aversion an individual with $T - t$ remaining periods of life consumes a fraction λ_{T-t} of her wealth, and this fraction is increasing in remaining lifetime (Laibson (1996)). Since individuals might benefit from a lottery win for a long time, it will be useful to consider T large, and we use the limit

at $T \rightarrow \infty$ to capture infinitely-lived individuals. This has no restriction for exponential discounters and naives. For sophisticates this constitutes a particular selection among all possible markov equilibria in infinite settings. It implies that individuals consume fraction λ_∞ of wealth, and this fraction increases in an individual's present-bias all else equal (see equation (9) in Laibson 1994). There exist other markov equilibria with different constant savings rate in these infinite games, and our results apply as long as they inherit the same comparative statics:

Assumption: Infinitely-lived hyperbolic discounters have Bernoulli utility function has constant absolute risk aversion $\rho \geq 1$, play an equilibrium markov strategy in the savings game where consumption is a constant fraction λ of wealth, and λ is increasing in impatience (β) all else equal.

A.5 Proof of Proposition 2

Consider infinitely-lived person i who wins the lottery reward R at time τ and has not other wealth. Consider her sequence of consumption choices C_0, C_1, C_2, \dots going forward, i.e., in periods $\tau, \tau + 1, \tau + 2$ etc. Now consider the same individual who wins the lottery reward R at time $t > \tau$ and has no other wealth. If $\tau > 0$, it is obvious that this person will choose exactly the same consumption sequence going forward: She will choose C_0, C_1, C_2, \dots in periods $t, t + 1, t + 2$ etc. The reason is that the environment going forward is exactly identical. That also implies that their continuation utilities are identical, so that $U_{i,\tau}(R) = U_{i,t}(R)$. This immediately establishes relationship (6) as a direct consequence of (5). This also holds for exponential discounters when $\tau = 0$ by the same logic.

It does not hold for hyperbolic discounters at $\tau = 0$. Consider first a sophisticate. At any point in time, this person is aware that her future selves are as present-biased as she is currently. Given our assumption on Markov equilibria (for which the limit of the game of finitely-lived players is a special case) the consumption sequence from time τ onwards rests exactly the same as the consumption sequence from time t onwards. So that step from the previous

paragraph remains. But the discount factors that are applied differ when the sum starts at zero compared to a future date:

$$U_{i,0}(R) = C_0 + \beta \sum_{s \geq 1} \delta^s u(C_s) < C_0 + \sum_{s \geq 1} \delta^s u(C_s) \leq U_{i,t}(R) \quad (12)$$

where in fact the last inequality holds with equality. By (probability-ratio-savings) it holds that $\varepsilon_\tau / \varepsilon_{i,t}^* \approx \gamma_{i,t,\tau} U_{i,t}(R) / U_{i,0}(R)$, so the probability ratio understates the true discount factor as stated in the proposition.

To make the same statement for naive hyperbolics we will use a related but slightly more sophisticated argument: let C_0, C_1, C_2, \dots be the optimal consumption sequence of a naive hyperbolic after lottery with wealth R at time zero. Then $U_{i,0}(R)$ can be constructed with the same equality as in (12). Also the strict inequality in (12) still holds. But now the weak inequality in fact holds strictly: from time t onwards the person could use the same consumption choices, but in fact he might even find a better sequence of consumption choices moving forward. Note that naives believe that their future selves will carry out their optimal decisions, so this logic applies. Again we conclude that the probability ratio understates the true discount factor as stated in the proposition.

We are left to show that nevertheless the probability ratio ranks individuals correctly. Consider two naive hyperbolic discounters i and j who only differ in respect to their present-bias parameter $\beta_i > \beta_j$ (i.e., they only differ in $\gamma_{i,0,t} > \gamma_{j,0,t}$ for all t). We have to show that $\beta_i \delta^t U_{i,0}(R) / U_{i,t}(R)$ is larger than $\beta_j \delta^t U_{j,0}(R) / U_{j,t}(R)$, i.e., that the probability ratio as in (5) is higher for person i than for person j . Here we omit the person index on the long-run discount factor because it is identical among them.

Clearly $U_{i,t}(R) = U_{j,t}(R)$ because in the future ($t > 0$) they expect both to discount exponentially with identical long-run discount factor. So we have to show that $U_{i,0}(R) / \beta_i$ is smaller than $U_{j,0}(R) / \beta_j$. Analogous to the previous arguments, consider a sequence of consumption choices C_0, C_1, C_2, \dots that is optimal for individual i at $\tau = 0$. Clearly: $\frac{U_{i,0}(R)}{\beta_i} = \frac{u(C_0)}{\beta_i} + \sum_{s>0} \delta^s u(C_s) < \frac{u(C_0)}{\beta_j} + \sum_{s>0} \delta^s u(C_s) \leq \frac{U_{j,0}(R)}{\beta_j}$, where the weak inequality arises because individ-

ual j might choose an even better sequence. This establishes the result for naives.

For sophisticates, we cannot use the same argument as they play a game rather than face an optimally chosen sequence, so in particular the last inequality of the previous argument is not obvious. So here we exploit brute-force the closed-form expression $\gamma_{i,0,t}U_{0,\tau}(R)/U_{i,t}(R) = \beta_i\delta_i^t [1 - (1 - \beta_i)\delta(1 - \iota)^{1-\rho}(1 - \lambda_i)^{(1-\rho)}]^{-1}$ when individuals save at rate λ_i , where we suppressed the person index on the right hand side for notational simplicity (see Laibson (1996), equation (29) for $U_{i,0}$, and for $U_{i,t}$ use the same equation but omit the present-bias). We will simply take comparative statics with respect to β_i directly, and indirectly through the change in λ_i . Clearly the first is positive. For the second, since λ_i is increasing in β_i , we have to show that the expression $\gamma_{i,\tau,t}\Upsilon_{i,\tau,t}$ is increasing in λ_i . This holds if $(1 - \lambda)^{1-\rho}$ is increasing in λ , or equivalently if $\rho \geq 1$. This completes the proof of Proposition 2.

A.6 Proof of Proposition 3

We want to show the following: Consider an infinitely-lived naive hyperbolic discounters with discount factors $\gamma_{i,\tau,t} = \delta_i^{t-\tau}$ if $\tau > 0$ and $\gamma_{i,0,t} = \beta_i\delta_i^t$ otherwise, who has a time-varying income stream $y_{i,t}$. She can save at person-specific interest rate τ_i^L in any period after winning our lottery and at rate τ_i otherwise. Fix any distance d , and fix a different infinitely lived naive hyperbolic discounter j . For R sufficiently large, there exists an open ball of winning probability around zero ε_τ such that for $\tau > 0$: $\gamma_{k,\tau,t} - \varepsilon_\tau/\varepsilon_{k,t}^* < d$ for each individual $k \in i, j$. Moreover, if both individuals are otherwise identical except for their short-run discount factors $\gamma_{i,0,t}$ and $\gamma_{j,0,t}$ and their usual interest rates τ_i and τ_j , then the probability ratio ranks correctly also relative to their short-run discount factor (i.e., if $\gamma_{i,0,t} < \gamma_{j,0,t}$ then $\varepsilon_\tau/\varepsilon_{i,t}^* < \varepsilon_\tau/\varepsilon_{j,t}^*$).

To show this, write agent $k \in \{i, j\}$'s problem in period $t > 0$ with net

present value W in the absence of our lottery as:

$$\begin{aligned} U_{k,+}(W, \iota) &= \max_c u_k(c) + \delta_k U_{k,+}(W') \\ \text{s.t.} \quad & W' = (1 + \iota)(W - c) \\ & W' \geq 0. \end{aligned}$$

Note that it is independent of the exact time period t . In case $t = 0$ it is

$$\begin{aligned} U_{k,0}(W, \iota) &= \max_c u_k(c) + \beta_k \delta_k U_{k,+}(W') \\ \text{s.t.} \quad & W' = (1 + \iota)(W - c) \\ & W' \geq 0. \end{aligned}$$

Let $W_k = \sum_s y_{k,s}/(1 + \iota_k)^s$ be the net present value of person k 's income stream when discounted at rate ι_k . From this initial net present value, denote by $C_{k,0}, C_{k,1}, C_{k,2}, \dots$ the sequence of consumption choices that maximize this recursive program. Standard arguments for such a simple recursive problem establish that $U_{k,+}(\cdot)$ and $U_{k,0}(\cdot)$ are strictly increasing, concave and differentiable. For ease of exposition write with slight abuse of notation $U_{k,t}(\cdot) := U_{k,+}(\cdot)$ when $t > 0$, even though the only variation in the continuation utility arises relative to time zero.

The ex-ante problem at time zero with a lottery that additionally pays R with probability ϵ at time $t > 0$ is then

$$\begin{aligned} \max_{c_0, c_1, \dots, c_{t-1}} \quad & u(c_0) + \beta_i \sum_{s=1}^t \delta_k^s u_k(c_s) + \beta_i \delta_k^t [(1 - \epsilon)U_{k,t}(W', \iota_k) + \epsilon U_{k,t}(W' + R, \iota_k^L)] \\ \text{s.t.} \quad & W' = W_k(1 + \iota_k)^t - \sum_{s=0}^t c_s(1 + \iota_k)^{t-s} \\ & W' \geq 0, \end{aligned}$$

Call the value of this program $\mathcal{U}_k(W_k, R, \epsilon, t)$. Clearly $C_{k,0}, C_{k,1}, \dots, C_{k,t-1}$ defined above are maximizers of this program when $\epsilon = 0$, and we can write $W'_{k,t} = W_k(1 + \iota_k)^t - \sum_{s=0}^t C_{k,s}(1 + \iota_k)^{t-s}$ for the net present value from period

t onward given this consumption path. If the lottery is already at time zero we have simply $\mathcal{U}_k(W_k, R, \epsilon, 0) = (1 - \epsilon)U_{k,0}(W_k, \iota_k) + \epsilon U_{k,0}(W_k + R, \iota_k^L)$.

The envelop theorem (e.g., Theorem 7, Morand et al. (2015))²¹ establishes:

$$\left. \frac{d\mathcal{U}_k(W_k, R, \epsilon, t)}{d\epsilon} \right|_{\epsilon=0} = U_{k,t}(W'_{k,t} + R, \iota_k^L) - U_{k,t}(W'_{k,t}, \iota_k).$$

That is, the (right-)derivative with respect to the winning probability can be computed as if the actual choices of consumption are unchanged. Therefore, we can write as first order Taylor approximation

$$\mathcal{U}_k(W_k, R, \epsilon, t) = \mathcal{U}_k(W_k, R, 0, t) + \epsilon[U_{k,t}(W'_{k,t} + R, \iota_k^L) - U_{k,t}(W'_{k,t}, \iota_k)] + O(\epsilon^2),$$

where $O(\epsilon^2)$ is the Bachmann–Landau notation for a function that vanishes at least at quadratic order when epsilon tends to zero, and $W'_{k,t}$ is the continuation net present value under the original consumption plan as defined above.

Recall that our elicitation method fixes an early winning probability ϵ_τ at time τ and elicits the winning probability $\epsilon_{k,t}^*$ at time t that makes person k indifferent, i.e., such that $U_k(W_k, R, \epsilon_\tau, \tau) = U_k(W_k, R, \epsilon_{k,t}^*, t)$. By the previous argument, for ϵ_τ close to zero this implies that $\epsilon_{k,t}^*$ has to be close to zero, and by the previous approximation this indifference be written as

$$\epsilon_\tau[U_{k,\tau}(W'_{k,\tau} + R, \iota_k^L) - U_{k,\tau}(W'_{k,\tau}, \iota_k)] \approx \epsilon_{k,t}^*[U_{k,t}(W'_{k,t} + R, \iota_k^L) - U_{k,t}(W'_{k,t}, \iota_k)].$$

This means that we can use the indifference condition approximately as if the person had a fixed consumption stream $C_{k,0}, C_{k,1}, \dots$ in the absence of a lottery win when the winning probability is small. All remaining arguments proceed along the lines of the proof for Proposition 2. This concludes the proof of Proposition 3.

²¹This particular envelop theorem is designed to accommodate parameters at the boundary of the domain; in our case: evaluation of the derivative at $\epsilon = 0$.

B Online Appendix: Additional Material on the Validation Experiment

B.1 Additional figures/tables

Table 8: Regression MLL estimates on coarse CBS estimates

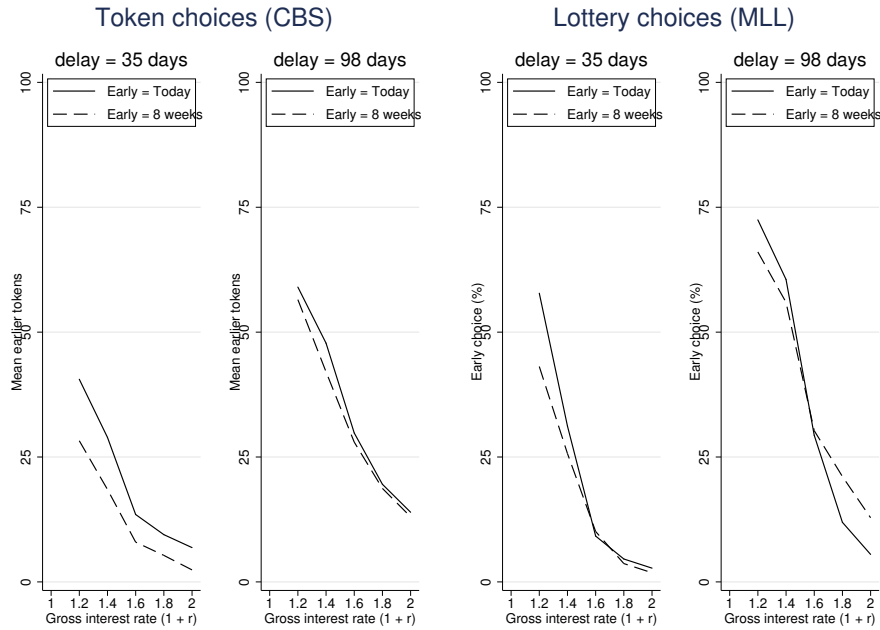
	Wave 1 (1) Delta (MLL)	Wave 2 (2) Delta (MLL)	Wave 1 (3) Beta (MLL)	Wave 2 (4) Beta (MLL)
Delta (CBS, coarse)	0.383* (0.221)	0.901*** (0.333)		
Beta (CBS, coarse)			0.191 (0.138)	0.388** (0.179)
Constant	0.476** (0.198)	-0.043 (0.299)	0.789*** (0.135)	0.630*** (0.173)
N	108	105	104	102

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 2: Average responses

(a) Wave 1 (Christmas)



(b) Wave 2 (Spring)

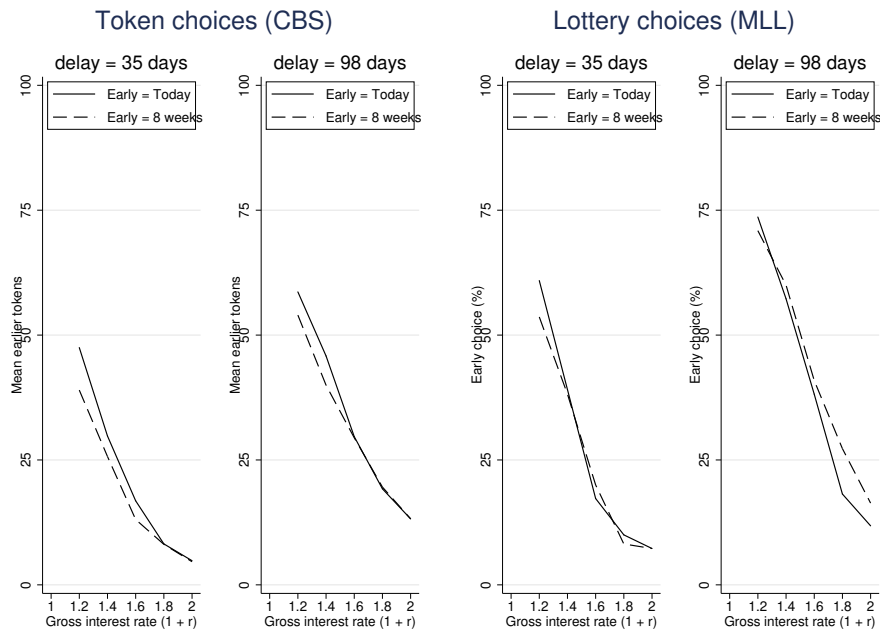


Figure 3: Wave 1: Distributions of preference parameters estimated at the individual level




Table 9: Regression MLL estimates on coarse CBS estimates by income/expenditure shocks (Delta)

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Delta (CBS, coarse)	0.991** (0.407)	0.168 (0.260)	1.494** (0.586)	0.641 (0.416)
Constant	-0.086 (0.362)	0.672*** (0.233)	-0.581 (0.529)	0.191 (0.370)
N	17	89	50	55

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 4: Wave 2: Distributions of preference parameters estimated at the individual level




Table 10: Regression MLL estimates on coarse CBS estimates by income/expenditure shocks (Beta)

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Beta (CBS, coarse)	0.210 (0.161)	0.165 (0.171)	0.696** (0.290)	0.103 (0.212)
Constant	0.766*** (0.158)	0.815*** (0.167)	0.378 (0.280)	0.864*** (0.203)
N	17	85	47	55

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.2 Details on Hypothesis 2

If the only difference between Wave 1 and Wave 2 is the Christmas shock, we expected the following.

Hypothesis 2: We expect a stronger present bias (lower β) in Wave 1 compared to Wave 2 when measured via CBS. We do not expect this measure to differ across both samples when measured via MLL.

As discussed in the main body, unfortunately differences across observables (age, gender and country of origin, see Table 2) make it unlikely that the premise is satisfied. To deal with these compositional differences, we present tests of Hypothesis 2 with and without controlling for these observed differences. We can, however, not exclude that in addition to the three observed variables, substantial differences exist on unobserved dimensions. Such differences render conclusions regarding Hypothesis 2 fairly unreliable.

We perform a one-sided t-test for equality of β between Wave 1 and Wave 2, both for the CBS and the MLL estimates. For CBS we do find that the mean β estimate is slightly lower in Wave 1 (0.988) than in Wave 2 (0.997), but the difference is not statistically significant. For MLL we find a difference in the same direction: 0.967 in Wave 1 and 1.002 in Wave 2 (with one-sided p-value of 0.087)

Alternatively we can count the share of present-biased classified individuals (those with $\hat{\beta} < 1$). The CBS method (coarse) yields 28% (Wave 1) and 30% (Wave 2), and a one-sided test for more present bias in Wave 1 fails to reject the null (p-value = 0.40). MLL yields 39% (Wave 1) and 31% (Wave 2) present-biased individuals and also here a one-sided test does not reject equality (p-value = 0.12). Thus, counting present-biased individuals fails to identify significant differences between waves, both for the CBS and for MLL methodology.

These results are likely confounded, at least partly, by the differences in samples between Wave 1 and 2. We proceed by performing the same tests, controlling for observed differences in age, gender and country of origin. To

those means, we estimate the following equation, where we pool all observation from the two waves together:

$$\hat{\delta}_i^{CBS} = \gamma_0 + \gamma_1 I_{\text{wave2}} + \beta X_i + \varepsilon_i \quad (13)$$

And identical for the MLL estimates. Results are shown in Table 11, where the model is estimated with and without control variables. First, we find some evidence that individual characteristics are predictive of the estimated preference parameters, although most coefficients are not statistically significant. Controlling for these characteristics has some impact on the differences between waves. For MLL (columns 3 and 4, first row), there is a small difference, which becomes slightly larger when adding controls. Using a one-sided test, these differences are statistically significant (at the 10% level). For CBS (columns 1 and 2), the wave-difference is not significant, but it does change when adding controls: after adding controls the difference is quite similar to the difference in the MLL results. We conclude that it seems likely that unobserved differences exist as well, rendering conclusions regarding hypothesis 2 less reliable.

Table 11: Regressions estimates with controls

	CBS ($\hat{\beta}^{CBS}$)		MLL ($\hat{\beta}^{MLL}$)	
	(1)	(2)	(3)	(4)
Wave 2	0.0086 (0.032)	0.032 (0.034)	0.035 (0.026)	0.043 (0.028)
Age 20-21		-0.065 (0.040)		-0.026 (0.033)
Age 22-30		-0.084 (0.071)		-0.10* (0.059)
Male		-0.051 (0.036)		-0.011 (0.029)
Netherlands		-0.036 (0.093)		-0.048 (0.071)
European		-0.037 (0.12)		-0.12 (0.091)
Log(Income)		-0.016 (0.020)		-0.0080 (0.016)
Constant	0.99*** (0.023)	1.17*** (0.16)	0.97*** (0.018)	1.08*** (0.13)
P-val 1-sided test Wave 2	0.39	0.17	0.087	0.059
N	216	213	216	213
r2	0.00034	0.039	0.0086	0.040

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

C Online Appendix: Parameter Estimation and Additional Empirical Robustness Analysis

C.1 Parameter estimation: CBS and MLL

Andreoni and Sprenger (2012) show that time preference parameters are related to the token choices in CBS questions as follows:

$$\ln\left(\frac{c_t - \omega_1}{c_{t+k} - \omega_2}\right) = \frac{\ln \beta}{\alpha - 1} \cdot 1_{t=0} + \frac{\ln \delta}{\alpha - 1} \cdot k + \frac{1}{\alpha - 1} \cdot \ln(1 + r)$$

Where c_t and c_{t+k} are chosen amounts in the early period t and the future period $t + k$, respectively. The additional utility parameters ω_1 and ω_2 can be interpreted as background consumption. Running a simple regression:

$$\ln\left(\frac{c_t - \omega_1}{c_{t+k} - \omega_2}\right) = \gamma_1 D_{t=0} + \gamma_2 k + \gamma_3 \ln(1 + r) \quad (14)$$

we obtain estimates that can be translated into parameters of interest (δ, α, β) :

$$\begin{aligned} \hat{\alpha} &= \frac{1}{\hat{\gamma}_3} + 1 \\ \hat{\delta} &= \exp\left(\frac{\hat{\gamma}_2}{\hat{\gamma}_3}\right) \\ \hat{\beta} &= \exp\left(\frac{\hat{\gamma}_1}{\hat{\gamma}_3}\right) \end{aligned} \quad (15)$$

Since we observe c_t and c_{t+k} , but not ω_1 and ω_2 , we opt for setting them equal to the monthly income as self-reported by participants in the experiment. We follow AS and estimate (14) using two-sided Tobit to take the bounds for the left-hand-side into account.²² This requires at least two interior choices per individual. For more than half of the sample this works, while we use OLS for about one-third (as AS do in the analysis in their appendix Table

²²The bounds are set at the individual level, with the lower bound being $\ln\left(\frac{0 - \omega_1}{c_{t+k} - \omega_2}\right)$ and the upper bound being $\ln\left(\frac{c_t - \omega_1}{0 - \omega_2}\right)$

A4). Almost all of the remaining individuals put 100% of their tokens on the patient chose, making a point estimate of the discount factor impossible. Rather than excluding them we include them and assign $\delta = \beta = 1$, which seems reasonable given their extreme patience.

The MLL questions are equivalent to the multiple price list methodology, in the sense that the switching point identifies an individual’s discount rate (or a discount rate *range*). AS define the discount rate by

$$\left(\frac{X}{Y}\right)^{\frac{1}{k}}$$

where X is the early amount (always 5), Y the later-payment at the switching point (6-10) and k is the interval (days between early and later). This approach implies that within the identified interval for the discount factor, the most “patient” value is picked. We pick the mid-point of the interval instead (both approaches are common in the literature). This only works for individuals with a unique switching point, but our sample is remarkably consistent: only in 2 out of 436 blocks in Wave 1, and 3 out of 403 blocks in Wave 2 there are multiple switching points (these participants are excluded). The long-term discount factor δ is identified from the long-run choices (where *early* is *8 weeks from now*), while β follows from the ratio between future discounting and immediate discounting. Histograms of the estimated preference parameters are presented in Figures 3 and 4 in the main paper (left panels for CBS and right panels for MLL).

C.2 Robustness analysis including 14-week interval questions for MLL estimation

In our baseline results we estimate preference parameters using only the 5-week interval MLL questions. That is, our estimate for the long-run discount factor δ follows from the switching point in the 8-week vs 13-week questions, while our short-run discount factor follows from the immediate vs 5-week questions. β is the ratio between the two. This approach is standard in the literature and requires few specification choices.

Since we also included questions with 14-week delays (8-weeks vs 22-weeks and immediate vs 14-weeks), we can essentially obtain two measures for the long-run and short-run discount factors. Since each measure identifies an interval rather than a point estimate, the overlap in the two interval might provide a more precise estimate. On the other hand, the two measures might produce non-overlapping intervals.

As a robustness check we reproduce our main results using MLL estimates based on *midpoint of the overlap in identified intervals*. In case of no overlap, we pick the midpoint between the intervals.

Results are presented in Table 12

Table 12: Regression MLL on CBS: Robustness analysis using MLL estimates from two intervals

	Wave 1 (1) Delta (MLL)	Wave 2 (2) Delta (MLL)	Wave 1 (3) Beta (MLL)	Wave 2 (4) Beta (MLL)
Delta (CBS)	0.323*** (0.120)	0.772*** (0.111)		
Beta (CBS)			0.123 (0.108)	0.169 (0.104)
Constant	0.559*** (0.111)	0.127 (0.102)	0.892*** (0.109)	0.899*** (0.107)
N	107	100	107	100

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 13: Robustness analysis: Regression MLL estimates on CBS estimates by income/expenditure shocks (Delta)

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Delta (CBS)	1.108*** (0.352)	0.204 (0.130)	0.979*** (0.160)	0.640*** (0.155)
Constant	-0.146 (0.311)	0.671*** (0.120)	-0.068 (0.150)	0.250* (0.142)
N	17	88	48	52

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 14: Robustness analysis: Regression MLL estimates on CBS estimates by income/expenditure shocks (Beta)

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Beta (CBS)	0.173 (0.186)	0.113 (0.126)	0.095 (0.116)	0.605** (0.283)
Constant	0.837*** (0.194)	0.903*** (0.126)	1.000*** (0.127)	0.458 (0.274)
N	17	88	48	52

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

C.3 Robustness analysis using ranks

As an alternative to regressing $\hat{\delta}^{MLL}$ on $\hat{\delta}^{CBS}$, we could instead compute the relative rank of the estimates *within Wave and method* and use these in regression. We transform the ranks to percentiles (ranging 0 to 1) and assign average ranks to ties (which occur a lot in MLL). Below are the regression results from a rank-rank regression.

They support the baseline results: coefficients are generally larger and statistically significant in Wave 2 and for individuals without financial shocks. Coefficients for β are smaller in general (as is the case in the baseline).

Table 15: Regression MLL on CBS

	Wave 1 (1)	Wave 2 (2)	Wave 1 (3)	Wave 2 (4)
	Rank(delta, MLL)	Rank(delta, MLL)	Rank(beta, MLL)	Rank(beta, ML
Rank(delta, CBS)	0.214** (0.085)	0.361*** (0.088)		
Rank(beta, CBS)			0.162* (0.091)	0.299*** (0.089)
Constant	0.393*** (0.049)	0.325*** (0.050)	0.419*** (0.053)	0.356*** (0.051)
N	108	105	108	105

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 16: Delta associations by financial shocks

	Wave 1		Wave 2	
	(1)	(2)	(3)	(4)
	No Shock	Shock	No Shock	Shock
Rank(delta, CBS)	0.788*** (0.216)	0.139 (0.093)	0.561*** (0.131)	0.230* (0.118)
Constant	0.160 (0.092)	0.438*** (0.056)	0.213*** (0.077)	0.392*** (0.066)
N	17	89	50	55

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 17: Beta associations by financial shocks

	Wave 1		Wave 2	
	(1) No Shock	(2) Shock	(3) No Shock	(4) Shock
Rank(beta, CBS)	0.090 (0.144)	0.159 (0.106)	0.363*** (0.116)	0.203 (0.137)
Constant	0.482*** (0.089)	0.414*** (0.060)	0.361*** (0.070)	0.367*** (0.074)
N	17	89	50	55

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$