

Heterogeneous risk and time preferences

Alina Ferecatu¹ · Ayse Önçüler²

Published online: 14 November 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract Assessing individuals' time and risk preferences is crucial in domains such as health-related decisions (e.g., dieting, addictions), environmentally-friendly practices, and saving opportunities. We propose a new method to jointly elicit and estimate risk attitudes and intertemporal choices. We use a novel individual level estimation procedure based on a hierarchical Bayes methodology, which can integrate different functional forms for discounting and risk attitudes. This method provides individual level estimates, and allows us to explore the heterogeneity in the data. In addition, we report a negative correlation between risk and time preferences, implying that risk-seeking individuals are less patient and less willing to defer consumption.

Keywords Risk and time · Risk preferences · Intertemporal choice · Joint elicitation · Hierarchical Bayes · Heterogeneity

JEL Classifications D81 · D90 · C11

Electronic supplementary material The online version of this article (doi:10.1007/s11166-016-9243-x) contains supplementary material, which is available to authorized users.

✉ Alina Ferecatu
ferecatu@rsm.nl

Ayşe Önçüler
onculer@essec.edu

¹ Rotterdam School of Management, Erasmus University, Burgemeester Oudlaan 50, 3062PA, Rotterdam, The Netherlands

² ESSEC Business School, 1 Avenue Bernard Hirsch, 95000, Cergy, France

1 Introduction

When making decisions in an intertemporal setting, individuals are known to heavily discount future outcomes and to have a strong preference for immediate gains (Loewenstein and Prelec 1992; Frederick et al. 2002). Understanding individuals' intertemporal decisions and designing tools to improve their decision making process has been a major research concern, with individual welfare and public policy implications. Are people myopic about health decisions such as following better diets, quitting smoking, or exercising? Do public policy makers respond appropriately to the threat of global warming, a decision involving a tradeoff between short term expenses and long run rewards? Chesson and Viscusi (2003) analyze the joint influence of time and uncertainty and conclude that people may have difficulties choosing the optimal precautionary measures to prevent climate change, a long term hazard, due to both ambiguity in the probability of global warming phenomena and also due to the ambiguity in the timing of global warming consequences. This suggests that a joint model of risk and time preferences is necessary to assess decision makers' tradeoffs between outcomes at different points in time.

The current study proposes a methodology that jointly elicits and estimates risk and time preferences at the individual level. We elicit risk attitudes and intertemporal choices following Holt and Laury's (2002) risk aversion experimental procedure and Collier and Williams's (1999) time discounting methodology involving price lists. We embed a constant relative risk aversion specification in an exponential discounting model, following Andersen et al. (2008). Our main contribution is methodological. We add a layer of flexibility to the estimation method through a hierarchical Bayes methodology, which allows us to recover individual level estimates and assess the heterogeneity in intertemporal choices and risk attitudes. Hierarchical Bayes modeling is a validated alternative to maximum likelihood estimation (Rossi et al. 2005; Toubia et al. 2013). The flexibility of the method in dealing with highly non-linear model specifications represents an important benefit. Such individual level estimates can be used in simulation studies to assess the effectiveness of changes in public policy decisions. To our knowledge, this is the first study that uses hierarchical Bayes modeling to jointly estimate risk and time preferences and to analyze the heterogeneity in the data.

Our results suggest that individuals are generally risk-averse, with a mean constant relative risk aversion coefficient of 0.515; their discount rates are in line with previous findings on joint elicitation (mean discount rate of 12%). The methodology is general and can be adapted to different functional forms. To test the robustness of our results, we estimate time preferences by using an exponential model, a hyperbolic model proposed by Mazur (1987), another hyperbolic specification by Prelec (2004), and a mixture model that allows a part of the population to behave as exponential discounters, and the remaining part as hyperbolic discounters. Our study provides evidence that controlling for the curvature of the risk aversion leads to lower estimated discount rates in experiments with hypothetical stakes. We conduct a clustering study and find three types of decision makers based on their risk and time preferences:

a high patience type (risk-averse, low discounting), a moderate patience type (with average risk and time preferences), and a low patience cluster (risk-seeking, high discounting). This result suggests that regulators could design policy interventions specifically aimed at people with various degrees of patience. The increased reliability of our estimates also allows us to revisit the issue of the correlation between risk and time preferences. We find evidence that risk-seeking people are more impatient and require higher interest rates to postpone consumption.

The remainder of the paper is organized as follows. In the next section, we present a brief literature review on joint elicitation and estimation of risk and time preferences. The experimental procedure for the joint elicitation of risk and time preferences is introduced in Section 3. In Section 4, we introduce our empirical model and present the estimation methodology involving a hierarchical Bayes approach. In Section 5, we report the results from our Bayesian estimation and further investigate the heterogeneity of risk and time preferences via individual differences and clusters of decision makers. The paper concludes with a discussion and the implications of our findings.

2 Existing literature

Recent developments in the intertemporal choice literature include a joint elicitation of risk and time preferences (Andersen et al. 2008; Laury et al. 2012; Andreoni and Sprenger 2012a; Toubia et al. 2013). This stream of literature emerged based on the conjecture that eliciting time preferences while assuming risk neutrality leads to an overestimation of discount rates. Using multiple price lists, Andersen et al. (2008) jointly elicit subjects' risk attitudes and intertemporal choices and show that discount rates are significantly lower than those reported in previous studies and more in line with market rates (10.1% on average). In another study, Andreoni and Sprenger (2012a) ask subjects to allocate a convex budget over a specific time period and find an average implied discount rate of 30% per year. They also observe that time preferences are generally dynamically consistent. Using a binary lottery mechanism, Laury et al. (2012) propose an elicitation of discount rates that corrects for the non-linearity, and is invariant to the form of the utility function. Their elicitation method yields similar results to the joint elicitation of time and risk preferences in Andersen et al. (2008), with an average discount rate of 12.2% per year. These results show that it is crucial to consider individuals' risk attitudes when eliciting time preferences in order to prevent an overestimation of discount rates.

The studies mentioned above as well as some other experimental findings (Kimball et al. 2009; von Gaudecker et al. 2011; Cardenas and Carpenter 2013) report significant heterogeneity in risk attitudes and intertemporal choices. Andersen et al. (2008) analyze the heterogeneity caused by observable individual characteristics as well as unobservable characteristics and find evidence for unobserved heterogeneity. Meyer (2013) reports substantive heterogeneity in time preferences for public goods, heterogeneity that cannot be explained by observable personal characteristics.

However, Chesson and Viscusi (2000) show that several individual difference variables impact time discounting. Nonsmokers and people with a post-graduate education degree have higher discount rates than smokers and subjects without advanced education. This highlights the challenge of designing a model of decision making under risk and time that can accommodate individual level behavior.

Several studies in behavioral decision making analyze the heterogeneity in the data and propose individual level parameter estimates. For instance, Jarnebrant et al. (2009) and Nilsson et al. (2011) estimate prospect theory parameters using hierarchical Bayes methods. A recent study by Toubia et al. (2013) focuses on building a dynamic (adaptive) survey design similar to conjoint studies, which presents decision makers with a personalized set of choices to elicit their risk and time preferences. They estimate cumulative prospect theory parameters as well as a quasi-hyperbolic time discounting model via a hierarchical Bayes procedure. They capture the heterogeneity in parameter values and use it for estimation and for designing the adaptive questionnaires.

Although essential for building theoretical and empirical models in decision making, as well as for public policy, the correlation between risk aversion and discount rate has received relatively little attention. Experimental studies on risk and time preferences have traditionally focused on one of the two phenomena apart from few exceptions. Andersen et al. (2008) report a small but significant positive correlation between subjects' attitudes toward risk and time, such that the level of discount rates increases with their risk aversion. Their findings are supported by evidence in Anderhub et al. (2001). Anderson and Stafford (2009) find that individuals become less patient as risk increases. Abdellaoui et al. (2013) report a small negative correlation between risk aversion and impatience for gains, in one of the two experiments they conducted. In a second experiment, they find no relation between risk aversion and impatience. Moreover, as utility over time and utility under risk are different, Abdellaoui et al. (2013) suggest using two utility functions, a risk function to transform risky prospects into certainty equivalents, and a time discounting function to recover the present value of the certainty equivalent. These contradictory results suggest that further assessment of the correlation between risk and time preferences is necessary.

3 Experimental design

We elicit risk and time preferences using the joint elicitation method proposed by Andersen et al. (2008).¹ The reason for choosing this methodology is twofold. First, we are interested in implied discount rates and our goal is to eliminate confounds introduced by the utility curvature that lead to upwards bias when estimating time

¹The experimental procedure can be found in the [Online Appendix](#), and is thoroughly documented in Andersen et al. (2008).

preferences. Second, the methodology is more pliable to our hierarchical Bayes specification, allowing us to obtain reliable individual level estimates of risk and time preferences.

We recruited subjects by sending an invitation email to MBA and M.Sc. students from a large French business school. Students were invited to attend the experiment during their lunch break, and were offered multiple time slots. While student samples are prevalent in decision making research, the literature raises issues related to the generalizability of the results to the overall population. Our subjects are enrolled in a business school and 96.5% of them have taken at least one finance related course prior to participating in the experiment. Thus, they seem to be more financially educated than the average population. Falk et al. (2013) study the behavior of student samples versus a representative sample of the general population and conclude that there is no concern for self-selection bias among students, and students' behavior does not seem to be significantly different than that of the general population. Our experimental results are in line with Laury et al. (2012), who employ an undergraduate student sample from a large US university. Andersen et al. (2008) use a representative sample of the adult Danish population and report an average discount rate lower than in our experiment. Andreoni and Sprenger (2012a) report a higher discount rate than all the above cited studies, with an undergraduate freshman and sophomore sample.

A total of 87 students participated in the experiment, in exchange for course credit. Subjects were informed that once the experiment is over, a lottery would be held, and 5 winners would receive a €20 gift voucher for a large media retailer. Subjects were not informed of the likely number of participants in the experiment, and were therefore not able to compute their chances of winning the lottery. The experiment began with the experimenter reading the instructions aloud. Participants were asked to turn to their computer screens and complete two tasks, the discount rate task and the risk aversion task. Subjects did four risk aversion tasks and four time discounting tasks, and completed a questionnaire about their age, gender, major, and their grade in the introductory finance course, if they had attended one. After completing the experimental task, subjects were debriefed and told that they would be informed about the results of the lottery within a month. Each experimental session took about 30 minutes. We conducted 21 experimental sessions, with about five subjects per session, to ensure sufficient attention from the experimenter to read the instructions and provide answers to questions.

3.1 Identifying risk preferences

The experimental procedure used to identify subjects' risk preferences is based on the Holt and Laury (2002) risk aversion experiment. Each subject is presented with a choice between two lotteries A and B. Table 1 shows the payoff matrix presented to the subjects. In the first decision line, lottery A offers a 10% chance of winning €200 and a 90% chance of winning €160. The respondents did not see the expected values of the lotteries or the constant relative risk aversion (CRRA) intervals. In Table 1, we can see that the expected values of lottery B increase relative to the expected values

Table 1 Typical payoff matrix in the risk aversion experiment. Subjects were asked to complete four risk aversion tasks with different payoffs, in the same income range

Lottery A		Lottery B				EV(A)	EV(B)	Difference	Open CRRA interval			
<i>p</i>	€	<i>p</i>	€	<i>p</i>	€	€	€	in €	if subject switches to lottery B			
0.1	200	0.9	160	0.1	385	0.9	10	164	47.5	116.5	−∞	−1.71
0.2	200	0.8	160	0.2	385	0.8	10	168	85	83	−1.71	−0.95
0.3	200	0.7	160	0.3	385	0.7	10	172	122.5	49.5	−0.95	−0.49
0.4	200	0.6	160	0.4	385	0.6	10	176	160	16	−0.49	−0.15
0.5	200	0.5	160	0.5	385	0.5	10	180	197.5	−17.5	−0.15	0.14
0.6	200	0.4	160	0.6	385	0.4	10	184	235	−51	0.14	0.41
0.7	200	0.3	160	0.7	385	0.3	10	188	272.5	−84.5	0.41	0.68
0.8	200	0.2	160	0.8	385	0.2	10	192	310	−118	0.68	0.97
0.9	200	0.1	160	0.9	385	0.1	10	196	347.5	−151.5	0.97	1.37
1	200	0	160	1	385	0	10	200	385	−185	1.37	∞

of lottery A, as one goes down the table. If a subject is indifferent between the two options, she can specify her indifference by choosing option I. The logic behind the risk aversion task implies that only very risk-seeking individuals will choose lottery B in the first row and only very risk adverse individuals will choose lottery A in the row before last. Risk neutral subjects are expected to switch to lottery B at decision row 5. The last row is a consistency check, as all subjects are bound to prefer lottery B, which strictly dominates lottery A.

For the sake of comparison, the relative payoffs are kept the same as in Andersen et al. (2008), but we divide the payoffs by 10 to remain in the income range over which we intend to estimate subjects' risk aversion.

3.2 Identifying time preferences

Our experimental design for eliciting time preferences is based on the multiple price lists introduced by Coller and Williams (1999). We follow similar time intervals as in Andersen et al. (2008), to be able to compare the two methodologies. We use four intertemporal choice tasks by manipulating the time when the rewards are offered to subjects ($t=2$ months, 6 months, 1 year and 2 years for option B). Table 2 is an example of the payoff table given to the subjects in order to elicit their time preferences. Subjects are asked to choose between option A and option B. If a subject is indifferent, she can choose option I, for indifference. Option A offers €300 in one month and option B offers €(300 + x) six months from now, where x is computed given a discount rate of 5% to 50% on the principal of €300, compounded quarterly. We use a compounded interest rate to increase ecological validity; in practice, a year or a quarter is the most common compounding period. We use the latter because in the first two

Table 2 Payoffs table for the 6-month time horizon in the discount rate experiment. Subjects completed four intertemporal choice tasks, with different time horizons (t=2 months, 6 months, 1 year and 2 years for option B)

Alternative	Payment Option A (amount (€) to be paid in 1 month)	Payment Option B (amount (€) to be paid in 6 months)	Annual Interest Rate (AR)	Annual Effective Interest Rate (AER)	Preferred Payment Option (Choose A, I or B)
1	300	306.17	5%	5.09%	A I B
2	300	312.40	10%	10.38%	A I B
3	300	318.68	15%	15.87%	A I B
4	300	324.99	20%	21.55%	A I B
5	300	331.36	25%	27.44%	A I B
6	300	337.78	30%	33.55%	A I B
7	300	344.25	35%	39.87%	A I B
8	300	350.76	40%	46.41%	A I B
9	300	357.32	45%	53.18%	A I B
10	300	363.92	50%	60.18%	A I B

time discounting tasks, the experimental time frames are less than a year. The elicited discount rates lie within a specific interval. For instance, for the choices available in Table 2, if a subject chooses option A in decision row 3 and option B in decision row 4, her discount rate lies in the interval (15%, 20%). Given that subjects complete four intertemporal choice tasks, our estimation method can infer stable interest rates that may lie inside an interval, not only at the interval bounds. We present subjects with annual and annual effective interest rates,² to facilitate comparison with other opportunities for investment present outside the lab and minimize errors in judgment.

We introduce a one-month delay for option A for all tasks to avoid quasi-hyperbolic discounting (Laibson 1997). There is extensive evidence in the economics and psychology literatures that individuals exhibit a bias for immediate payoffs. Discount rates elicited from choices between a current payoff and a future one are significantly higher than the discount rates elicited from two future choices. One explanation for such present biased preferences is a preference for certain outcomes (Andreoni and Sprenger 2012b). Subjects would disproportionately prefer the certain present option to the inherently risky future option. Since this present bias is not the focus of our study, we introduced a front-end delay. In our experiment, subjects choose between two future payoffs and they are expected to behave as exponential discounters (Andreoni and Sprenger 2012b). Delaying both options also reduces the influence of transaction costs on intertemporal choices. If only the delayed option

²These terms are explained in the experimental instructions (see the [Online Appendix](#)).

involved greater transaction costs, the revealed discount rate would include these subjective costs. As both options occur in the future, such transaction costs are equivalent and will not influence the revealed discount rate.³

3.3 Consistency and validity checks

Our experiment involves hypothetical stakes. The literature presents mixed evidence on whether individuals respond differently to hypothetical choices involving risk or time discounting than they do in a real-choice context. Taylor (2013) shows that on average, measured risk preferences are not significantly different across real and hypothetical settings. Holt and Laury (2002) demonstrate that individuals appear to be more tolerant of risk in a hypothetical setting, and that this “hypothetical bias” increases with the size of the stakes. Camerer (2004) notes that the effect of incentives on behavior is likely to depend on the task. However, the choice over money gambles is not likely to be a domain in which people would behave according to expected utility theory if they put more effort into the task. Using functional magnetic resonance imaging (fMRI), Kang et al. (2011) show that common areas of the brain are activated when individuals make real and hypothetical choices about the purchase of consumer goods, but that the level of this activity differs.

In our experiment, we provided the subjects with a sample task before the experiment to ensure that they understood the instructions clearly (one decision row was chosen from both time discounting and risk preference tasks). We also performed a consistency check for the first time discounting and risk aversion tasks (see the [Online Appendix](#)). Few subjects switched between lotteries A and B in either tasks. Subjects switched back to lottery A after choosing lottery B in only 1.7% of the risk aversion tasks (6 tasks out of 348 completed) and in 3.7% of the time discounting elicitation (13 tasks out of 348 completed). Furthermore, no subject chose only option A or only option B throughout the four time discounting or risk aversion tasks, suggesting that subjects’ risk and time preferences are within the ranges provided in the experimental design.

To assess the quality of our data, we conducted a within-subject analysis in the four risk aversion and time discounting tasks. After computing the variance in switching points, we averaged the standard deviation across subjects. On average, the switching point between the smaller, sooner and larger, later payoffs varied by 1.66 decision rows. For the risk preferences task, the switching point varied on average by 1.09 decision rows. The switching point varied significantly more for the discount rate tasks than for the risk aversion tasks ($p < 0.01$). This variation is not only due to response error, but also due to hyperbolic discounting, i.e. subjects tend to be more patient, thus switch at a lower decision row, as the time interval increases. Even though most subjects chose to switch from lottery A to lottery B without selecting the option corresponding to an indifference point (79.04% for risk aversion tasks, 69.5% for the discounting tasks), the relatively large number of stated indifference points

³We note that the two-month delay for option B in the first intertemporal choice task occurred during the summer vacation. This could have led to an increase in reported discount rates.

(20.96% for the risk elicitation task and 30.5% for the discounting task) helped us to estimate the discount rates and risk aversion coefficients, by improving the precision of the recovered estimates.

4 Model specification

Our hierarchical Bayes framework provides a natural context in which we jointly estimate individual level risk and time preferences. We set up 1) a risk preference specification, 2) a method to jointly measure risk and time preferences, and 3) a method to capture heterogeneity in individual preferences. We review all the three components in turn.

4.1 Modeling risk preferences

We do not impose any assumptions about linear risk preferences in income. Consequently, the utility function is the constant relative risk aversion (CRRA) specification (Holt and Laury 2002):

$$U(M) = \frac{(\omega + M)^{(1-r)}}{1 - r}, \text{ for } r \neq 1 \tag{1}$$

where M is the monetary amount, ω is the time invariant amount of background consumption, and r is the CRRA coefficient where $r = 0$ denotes risk neutral behavior, $r > 0$ denotes risk aversion and $r < 0$ denotes risk-seeking behavior. The background consumption represents estimated daily consumption levels, assumed constant over the time frame we analyze. Given our CRRA specification and positive values of background consumption, as opposed to $\omega=0$, observed choices from our experiment imply higher levels of risk coefficients, thus more curvature of the utility function. We assume positive levels of background consumption as in Andersen et al. (2008) and exogenously vary its value to provide a sensitivity analysis of our results to the level of estimated daily consumption.

To specify the likelihood of the risk aversion responses, we first compute the expected utility (EU) as the sum of payoffs weighted by the probabilities induced by the experimenter. Given that there are two outcomes ($k=1, 2$) in each lottery j , the EU for each lottery is:

$$EU_j = \sum_{k=1,2} p(M_k) \times U(\omega + M_k) \tag{2}$$

Following Holt and Laury (2002), we implement a simple stochastic specification to specify the likelihood conditional on the model. The specification is standard in the decision making literature and has been implemented by Andersen et al. (2008) and Laury et al. (2012).

$$\nabla EU = \frac{EU_B^{\frac{1}{\mu}}}{EU_B^{\frac{1}{\mu}} + EU_A^{\frac{1}{\mu}}} \tag{3}$$

where ∇EU is the cumulative probability distribution reflecting differences in the expected utility of the two lotteries A and B and μ is a behavioral noise parameter.

When $\mu \rightarrow 0$, the specification collapses to the deterministic expected utility theory (EUT), where the alternative with the highest expected utility will be chosen with certainty. However, when $\mu \rightarrow \infty$, the choices become random, indifferent to the expected utilities. Toubia et al. (2013) introduce a similar error specification, where a noise parameter implies whether choices are random or whether they converge to a deterministic selection.

The likelihood of the risk aversion responses conditional on EUT and CRRA specification depends on r and μ . The conditional log-likelihood is:

$$\ln L^{RA}(r, \mu, y, \omega) = \sum_j \{[\ln(\nabla EU)|_{y_j = 1}] + [\frac{1}{2}\ln(\nabla EU) + \frac{1}{2}\ln(1 - \nabla EU)|_{y_j = 0}] + [\ln(1 - \nabla EU)|_{y_j = -1}]\} \quad (4)$$

where $y_j = 1$ implies the choice of option B; $y_j = -1$ the choice of option A and $y_j = 0$ shows indifference between A and B, implying a 50-50 mixture of the likelihood of choosing either of the lotteries.

4.2 Modeling time preferences

We first consider the mainstream specification of time preferences. Assuming exponential discounting, a subject is indifferent between two income options if:

$$U(\omega + M_t) + \frac{1}{(1 + \delta)^\tau} U(\omega) = U(\omega) + \frac{1}{(1 + \delta)^\tau} U(\omega + M_{t+\tau}) \quad (5)$$

where $U(\omega + M_t)$ is the utility of monetary outcome M_t for delivery at time t plus background consumption ω ; δ is the discount rate; τ is the delay to the larger, later reward.

The utility function $U(\cdot)$ is separable and stationary over time. The left-hand side of Eq. 5 is the sum of the discounted utilities of receiving the monetary outcome M_t at time t (in addition to background consumption) and receiving nothing extra at time $t + \tau$, and the right-hand side is the sum of the discounted utilities of receiving nothing over background consumption at time t and the outcome $M_{t+\tau}$ (plus background consumption) at time $t + \tau$.

The specification we use for the discount rate is similar to the one employed for risk aversion. By assuming the CRRA utility function as before, we can write the discounted utilities of each of the two options as:

$$PV_A = \frac{1}{(1 + \delta)^t} U(\omega + M_A) + \frac{1}{(1 + \delta)^{t+\tau}} U(\omega) \quad (6)$$

$$PV_B = \frac{1}{(1 + \delta)^t} U(\omega) + \frac{1}{(1 + \delta)^{t+\tau}} U(\omega + M_B) \quad (7)$$

where PV_A and PV_B are the present values of option A and option B; M_A and M_B are the monetary amounts in the discounting tasks presented to subjects, to be paid at time t and time $t + \tau$ respectively.

The cumulative distribution function for the discount rate task is specified similarly as for risk aversion.

$$\nabla PV = \frac{PV_B^{\frac{1}{\nu}}}{PV_B^{\frac{1}{\nu}} + PV_A^{\frac{1}{\nu}}} \tag{8}$$

where ∇PV is the cumulative probability distribution reflecting differences in the present values of the two tasks, conditional on r , δ and ν . ν is the behavioral noise parameter, allowing some subjects to make some errors from the perspective of expected utility theory, similar to the μ noise parameter in the risk aversion task. Given that the risk aversion task involves choices in lotteries as opposed to values, we assume that it is cognitively more difficult, therefore expect $\mu > \nu$.

The log-likelihood for the time discounting responses, conditional on r , δ and ν , is:

$$\begin{aligned} \ln L^{DR}(r, \mu, \delta, \nu, y, \omega) = & \sum_j \{[\ln(\nabla PV)|y_j = 1] + [\frac{1}{2}\ln(\nabla PV) + \\ & + \frac{1}{2}\ln(1 - \nabla PV)|y_j = 0] + [\ln(1 - \nabla PV)|y_j = -1]\} \end{aligned} \tag{9}$$

where $y_j = -1, 1$ and 0 denotes selection of option A, B or indifference in observation j , respectively.

The joint likelihood of the risk aversion and the discount rate responses can then be written as:

$$\ln L(r, \mu, \delta, \nu, y, \omega) = \ln L^{RA} + \ln L^{DR} \tag{10}$$

4.3 The hierarchical Bayes specification: Heterogeneity and priors

The most common approach in estimating decision making models (cumulative prospect theory, time discounting etc.) is single-subject maximum likelihood estimation (MLE). There are several limitations of using a single-subject MLE methodology. First, while the method is efficient when a large number of observations per subject are available, its main drawback is the assumption that all subjects are independent, without considering that individual parameter estimates originate from a group-level distribution. The hierarchical Bayes estimation simultaneously estimates the individual level parameters and addresses the above limitation through a compromise between the two extremes of complete independence and pooling (average-level parameter estimates where all subjects are presumed to be identical). The method shrinks the individual estimates towards the group mean, and this effect is more pronounced when the individual estimates are less reliable.

Second, the single-subject MLE provides point estimates for the model parameters. Using these point estimates in analyses of variance or regressions, without factoring in their reliabilities, could lead to poor predictions, as results can be severely influenced by extreme observations. Through the above mentioned shrinkage effect,

the hierarchical Bayes estimation provides a way to directly quantify the uncertainty on the parameters. Thus, it prevents unreliable information from having a disproportionate influence on the parameter estimates.

A random effects approach (or an error components approach—formally equivalent) allows for correlation in unobserved utility over several choices for each respondent, and overcomes the assumption of complete independence. Since such a model usually involves high-dimensional integrals that are analytically intractable, the random effects model is estimated via maximum simulated likelihood (MSLE), or Bayesian methods. In addition to the simulation step necessary to integrate over all values of the model parameters for each individual, MSLE also involves an optimization step, leaving it vulnerable to issues such as reach of global maxima vs. local maxima, and requires a tolerance level and proper starting values. Bayesian inference is a simulation technique that can overcome such problems, because it covers the parameter space more effectively (see Jackman 2009 for proofs involving hierarchical specifications). The quantities of interest in Bayesian inferences are the posterior distributions of the model parameters, and the method provides an *exact* estimation of these distributions, in small as well as in large samples (Rossi and Allenby 2003). MSLE provides an *approximation* of these distributions, and needs large samples for asymptotic properties of convergence. Relevant for our setting, the Bayesian inference is more suited than MSLE to ensure that all parameters have the same sign or are bounded within a certain range for all subjects. The Bayesian procedure handles transformations of the model parameters more efficiently (e.g., exponential transformation for log-normally distributed parameters), as it does not search for a maximum. The method is particularly suited in situations where only a limited amount of data per subject are available for a moderately large sample of subjects, typical for experiments in decision making. It provides reliable estimates even with a small amount of data per individual (Rossi and Allenby 2003).

Using Bayes' theorem, we combine the likelihood function in Eq. 10 with prior distributions to obtain posterior distributions for our model parameters. The two building blocks of our Bayesian framework are related to the specification of heterogeneity and the introduction of prior distributions of the parameters. We introduce each of the aspects in turn.

We estimate individual level parameters by leveraging the information on the distribution of parameters across subjects in our experiment. We specify the following heterogeneity model:

$$\theta_i = \theta_0 + u_i; u_i \sim N(0, V_\theta) \quad (11)$$

where θ_i is the vector of individual level risk and time coefficients $(r_i, \delta_i, \mu_i, \nu_i)$. θ_0 represents a vector of average risk and time preferences and related noise parameters. The unexplained heterogeneity, u_i , is normally distributed. V_θ shows the variance of the model parameters and the covariances between them. This matrix, along with the correlation between the individual level risk and time estimates, allows us to assess the link between subjects' risk attitudes and intertemporal choices.

Our model is a hierarchical Bayes model because we formulate prior distributions not only on the parameters that specify the likelihood function, but also on those

priors themselves. We set up a two-stage model that allows the first-stage priors to be influenced by the data. The prior on θ_i , usually referred to as a first-stage prior, follows a normal distribution:

First-stage prior

$$\theta_i \sim N(\theta_0, V_\theta)$$

To ensure the parameters remain in an acceptable range, i.e., δ_i, μ_i, v_i are positive-definite, we apply an exponential transformation to the parameter estimates. We also set priors on the parameters in the heterogeneity model, called second-stage priors.

Second-stage prior

$$\begin{aligned} \text{vec}(\theta_0|V_\theta) &\sim N(\bar{\theta}_0, A^{-1} \otimes V_\theta) \\ V_\theta &\sim IW(v, V_0) \end{aligned}$$

The second-stage priors would allow us to incorporate any information that we might have before our data collection. These priors are usually chosen to be as uninformative as possible, to allow the data to determine the values of the parameter estimates. All the above specifications are standard in the hierarchical Bayes literature (Rossi et al. 2005). The prior distributions for θ_0 and V_θ are a normal distribution and an Inverse Wishart distribution, respectively. The latter is chosen because it is a conjugate prior to the likelihood implied by the first-stage prior, thus the posterior distribution of V_θ is also Inverse Wishart. We obtain the posterior distribution of the model parameters by combining the likelihood in Eq. 10 with the first and second-stage priors:

$$\begin{aligned} p(\{\theta_i\}, \theta_0, V_\theta|Data) &\propto p(Data|\{\theta_i\}) \times p(\{\theta_i\}|\theta_0, V_\theta) \\ &\times p(\{\theta_0\}|\bar{\theta}_0, A^{-1}) \times p(V_\theta|v, V_0) \end{aligned} \quad (12)$$

where $p(Data|\{\theta_i\})$ is given by the likelihood in Eq. 10, $p(\{\theta_i\}|\theta_0, V_\theta)$ is the first-stage prior, $p(\{\theta_0\}|\bar{\theta}_0, A^{-1})$ and $p(V_\theta|v, V_0)$ are the second-stage priors. The value of hierarchical Bayes lies in estimating the parameters for all individuals simultaneously, as opposed to traditional methods, which estimate the parameters of each decision maker independently. The model estimates parameters for each individual, but restricts these parameters through the group distribution. This constraint allows for the potentially unreliable information from one participant to be weighted against the information from all other individuals, and prevents overfitting.

We make inferences about all model parameters based on posterior distributions obtained through Markov Chain Monte Carlo (MCMC) methods. We use the point estimates of risk and time parameters for each subject in our experiment, drawn from the posterior distribution, to characterize the extent of heterogeneity.⁴

⁴Additional details on data preparation, model estimation, and the code are available upon request.

4.4 Alternative specifications

We test the exponential discounting model against alternative specifications. We first implement Mazur's (1987) hyperbolic specification and replace Eqs. 6 and 7 with:

$$PV_A = \frac{1}{(1 + \gamma t)} U(\omega + M_A) + \frac{1}{(1 + \gamma(t + \tau))} U(\omega) \quad (6')$$

and

$$PV_B = \frac{1}{(1 + \gamma t)} U(\omega) + \frac{1}{(1 + \gamma(t + \tau))} U(\omega + M_B) \quad (7')$$

where $\gamma > 0$ is the discount rate.

We also evaluate the effect of using a more general hyperbolic specification proposed by Prelec (2004). This specification replaces Eqs. 6 and 7 with:

$$PV_A = \exp(-\beta t^\alpha) U(\omega + M_A) + \exp(-\beta(t + \tau)^\alpha) U(\omega) \quad (6'')$$

and

$$PV_B = \exp(-\beta t^\alpha) U(\omega) + \exp(-\beta(t + \tau)^\alpha) U(\omega + M_B) \quad (7'')$$

where α exhibits decreasing impatience, i.e., lower discount rates over time. When $\alpha = 1$, the specification collapses to the exponential discounting model. β characterizes time preferences in the conventional sense. It represents the instantaneous discount rate when $\alpha = 1$. When $\alpha < 1$, the instantaneous discount rate is given by $\alpha\beta t^{\alpha-1}$.

We also consider a mix of populations behaving differently for choice under risk and over time (Conte et al. 2011; Andersen et al. 2008). From a discounting theory perspective, this implies that a proportion of the population discounts exponentially, while the rest are hyperbolic discounters. We consider the sensitivity of our results to a statistical specification allowing more than one latent process to generate each observation. Following Andersen et al. (2008), we introduce a finite mixture model of exponential and hyperbolic models of discounting. The mixture likelihood function is given by:

$$\ln L(r, \mu, \delta, v, y, \omega) = \ln L^{RA} + \pi \ln L^{DR-E} + (1 - \pi) \ln L^{DR-H} \quad (10')$$

where π is the probability that a given observation is generated by the exponential discounting model, such that $0 \leq \pi \leq 1$. $\ln L^{DR-E}$ represents exponential discounting and $\ln L^{DR-H}$ represents Mazur's or Prelec's hyperbolic discounting.

5 Experimental results

5.1 Descriptive statistics

We estimate individual level risk and time preferences using data from our experiment. For the estimation, we set the background consumption equal to $\omega = 12$ in

order to compare our results with previous literature (Andersen et al. 2008). Table 3 presents the descriptive statistics for the risk aversion and discounting experimental tasks. The risk aversion measure shows the average decision row where subjects switched from option A to option B. On average, subjects chose the “safe” lottery A for six rows. A risk neutral subject would choose Lottery A for the first four decision rows and then switch to Lottery B. Thus our average subject is moderately risk-averse. The average discount rate is about 26%. The Pearson correlation between risk aversion (the average decision row where subjects switch in the four experimental tasks) and the model-free discount rate is 0.14, not significantly different from 0 ($p = 0.17$).

Figure 1 displays the proportion of subjects selecting the sooner option (payment in one month) in each of the decision rows, for the four discounting tasks. Starting in decision row 3, a higher proportion of subjects choose the sooner option A over the later option B in the two-month delay task than in the two-year delay task. This implies that subjects are most impatient in the two-month delay task and least impatient in the two-year delay task. This is in line with previous literature and provides some model-free evidence of hyperbolic discounting.

5.2 Estimation results

We estimate the posterior distribution of the parameters based on a total of 10,000 MCMC draws, and assess convergence visually by inspecting the MCMC trace plot of the log-likelihood. The MCMC chain converged after 5,000 iterations, thus we discard the first 5,000 draws. Table 4 summarizes the posterior distributions of the risk and time parameters across individuals.

We use the log Bayes Factor (logBF) to compare models. This measure accounts for model fit and prevents overfitting (Kass and Raftery 1995). The measure integrates over the set of parameters in each model. Therefore, Bayes factors automatically penalize models with too much complexity (too many parameters). Table 4 shows all the model comparisons necessary to determine which data generating process fits the data better, i.e. exponential discounting, hyperbolic discounting or a mixture model. The logBF is the difference between the log-marginal likelihood (LML) of model 2 (M2) and model 1 (M1). We use the Newton Raftery estimator to estimate the log marginal likelihood (i.e. the harmonic mean of log-likelihood

Table 3 Descriptive statistics of the risk measures and the discount rates (n=87)

	Min	Max	Average	Std. Deviation
Discount Rate	0	50	25.91	13.52
Risk Aversion	2.25	9	5.68	1.61
Age	20	37	23.89	3.41
Gender	41 men	46 women	–	–

draws after the burn-in period). Kass and Raftery (1995) suggest that a value of $\log\text{BF}=\text{LML of M2} - \text{LML of M1}$ higher than 5 provides evidence of the superiority of model 2. Here, we benchmark all models against the exponential discounting model (M1).

5.2.1 Model comparison

The comparative results suggest that the model that best fits the data is Prelec's hyperbolic model. The value of the α parameter is significantly lower than 1 ($\alpha = 0.748$, $s.e. = 0.072$), suggesting that participants discount more heavily rewards received in the near future. Figure 2 compares the instantaneous discount rates for the hyperbolic and exponential models, assuming risk aversion and risk neutrality. Both the exponential and Prelec's hyperbolic model show the same basic results regarding elicited discount rates. Failing to correct for the curvature of the utility function when eliciting time preferences results in overestimating discount rates. We use both Mazur's hyperbolic and Prelec's hyperbolic model to investigate a mixture specification of exponential and hyperbolic discounting. Both specifications reveal similar results. About 60% of the subjects discount future rewards hyperbolically, while the remaining 40% discount exponentially, showing that individuals have heterogeneous underlying discounting patterns. Given the widespread usage of exponential discounting and for the sake of comparison with results from earlier studies, we provide a heterogeneity analysis for the exponential model in Section 5.3.

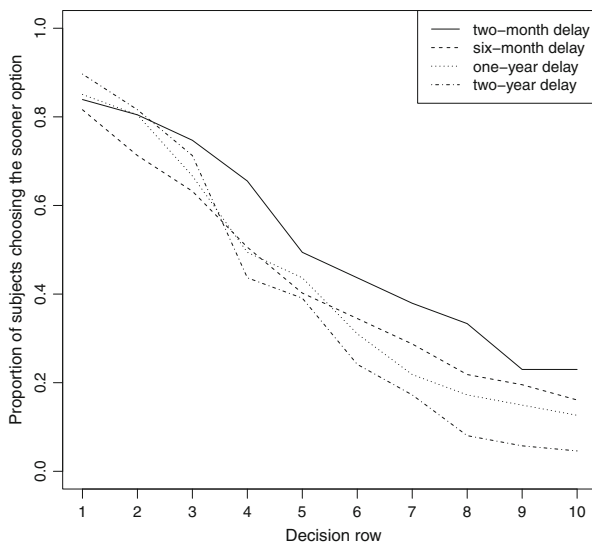


Fig. 1 Proportion of participants choosing the sooner option (option A) by decision row in the four discounting tasks

5.2.2 Time preferences

Our results suggest that, on average, individuals demand relatively high interest rates to delay consumption, and are moderately risk-averse. As depicted in Table 4, the discount rate parameter has a posterior mean of 12.05% with the 95% credible interval at (9.93%, 14.49%); the average constant relative risk aversion is 0.515, with the 95% credible interval at (0.418, 0.614). Assuming risk neutrality (the constant relative risk aversion coefficient is set to 0), the discount rate is estimated at 22.79% with a 95% credible interval at (19.19%, 26.84%). The average discount rates after correcting for the curvature of the utility function are in line and strengthen the results from the literature on joint elicitation of risk and time preferences (Andersen et al. 2008; Andreoni and Sprenger 2012a). Moreover, our estimated discount rates are much lower than what was previously reported in the intertemporal choice literature, where discount rates of the order of hundreds are not uncommon (Loewenstein and Prelec 1992; Frederick et al. 2002). Given our sample of financially educated business school students, we might expect discount rates to be lower than in other studies. Our estimated discount rates are indeed below the level reported in Andreoni and Sprenger (2012a), who used a freshman and sophomore sample in their study. The

Table 4 Parameter estimates: means and standard deviations (below) of the posterior distributions (n=87)

Model	Risk Aversion	Noise Parameter μ	Discount Rate	Noise Parameter ν	Mixture Parameter π	logBF ³ (LML)
Exponential	0.515 ^{1a} (0.05)	0.041 ^a (0.004)	0.120 ^a (0.011)	0.011 ^a (0.001)		– (–2,010)
Mazur’s hyperbolic	0.507 ^a (0.051)	0.043 ^a (0.004)	0.126 ^a (0.012)	0.009 ^a (0.001)		39 (–1,971)
Prelec’s hyperbolic (β, α)	0.514 ^a (0.052)	0.041 ^a (0.004)	0.131, 0.748 ^a (0.015, 0.072)	0.008 ^a (0.001)		210 (–1,800)
Exponential/Prelec’s hyperbolic Mixture (δ, β, α)	0.41 ^a (0.051)	0.040 ^a (0.004)	0.113 0.135, 0.651 ^a (0.021, 0.098)	0.008 (0.001)	0.4106 ² (0.063)	136 (–1,874)

¹Statistical signif. codes:(a) 99% interval does not include 0

²The mixture parameter π is not statistically different than 50%

³Last column reports the log-Bayes factors and log-marginal likelihoods (in brackets)

The logBF is the difference between the log-marginal likelihood of the specified model minus the log-marginal likelihood of the exponential discounting model. A value higher than 5 suggests that the specified model fits the data better

results are in line with Laury et al. (2012) findings, estimated with an undergraduate student sample from a large US university. Also, our results show slightly higher discount rates than reported in Andersen et al. (2008), which used a representative sample of the adult Danish population.

5.2.3 Risk preferences

Our utility curvature estimates are lower than those reported in Andersen et al. (2008) (average CRRA estimate of 0.515; s.e. = 0.050 versus 0.741; s.e. = 0.048). In turn, our results imply a higher level of risk aversion than those reported by Andreoni and Sprenger (2012a), who find that aggregate utility curvature is far closer to linear utility than estimated from the double multiple-price list approach used in Andersen et al. (2008). Interestingly, they find no correlation between the risk estimates from the multiple price list approach and the convex time budget approach, despite the fact that under CRRA utility, the two elicitation methods measure the same utility construct. This suggests that further research is necessary to assess which methodology is more suited to assess curvature in discounting models. The noise parameters, showing the extent to which subjects' choices are due to decision errors or randomness, are also within expected ranges (Holt and Laury 2002; Andersen et al. 2008). The fact that the noise parameter estimate is smaller for the time discounting task than for the risk aversion task (0.041; s.e. = 0.004 for risk aversion versus 0.011; s.e. = 0.001 for discounting) is not surprising either, as the




Fig. 2 Estimated instantaneous discount rate for hyperbolic discounters vs. exponential discounters, assuming risk aversion (left) and risk neutrality (right)

time discounting elicitation procedure is cognitively easier than the risk preference task.

5.3 Heterogeneity in risk and time preferences

5.3.1 Investigating possible drivers of heterogeneity—observable characteristics

The heterogeneity model presented in Section 4.3 is a special case of a hierarchical Bayes random effects model, in which the distribution of heterogeneity is recovered and described, but the heterogeneity in preferences is not explained by observed variables. We extend this model and attempt to partly explain heterogeneity in risk and time preferences using the subjects’ individual differences. We collected data on gender, age, and grade in the introductory finance course if they had attended one (84/87 or 96.5% of the sample). We conduct our analysis based on the exponential discounting model. Equation 11 is replaced with:

$$\theta_i = \Lambda z_i + u_i; u_i \sim N(0, V_\theta) \tag{11'}$$

where z_i is a set of individual difference variables for individual i , and Λ captures the impact of covariates on subjects’ risk and time preferences. The first and second-stage priors will also reflect the changes in our model.

First-stage prior

$$\theta_i \sim N(\Lambda z_i, V_\theta)$$

Second-stage prior

$$vec(\Lambda|V_\theta) \sim N(\bar{\Lambda}, A^{-1} \otimes V_\theta)$$

$$V_\theta \sim IW(\nu, V_0)$$

Table 5 Parameter estimates for the observed heterogeneity model, showing the impact of the individual difference variables on risk and time preferences: means and standard deviations (below) of the posterior distributions (n=87)

Model Parameters	Risk Aversion	Noise Parameter μ	Discount Rate	Noise Parameter ν
Intercept	0.487 ^{1a} 0.046	0.049 ^a 0.004	0.123 ^a 0.012	0.012 ^a 0.001
Age	0.002 0.014	-0.002 0.028	-0.065 ^a 0.028	-0.044 0.028
Gender	0.180 ^c 0.098	-0.110 0.180	-0.096 0.177	-0.448 0.186
Finance Grade	-0.002 0.021	-0.031 0.03f9	0.035 0.039	0.065 0.041

¹Statistical signif. codes: (a) 99% interval does not include 0. (c) 90% interval does not include 0

We estimated the model following the methodology outlined in Section 4. We mean-centered the demographic variables to be able to interpret the intercept as the average parameters.

The explained heterogeneity model converged to a log-marginal likelihood of $-2,098$. The logBF between the explained heterogeneity and the unexplained heterogeneity model is 88. Therefore, we find strong evidence that the unexplained heterogeneity model is more parsimonious and seems to fit the data better, as our individual difference variables do not seem to explain well the variance in subjects' risk attitudes and intertemporal choices. In line with previous literature, we find no significant effect of gender on intertemporal choices (Chesson and Viscusi 2000). Table 5 reports two significant effects. Gender positively impacts risk aversion ($\lambda_{r-gender} = 0.180$; $s.e. = 0.098$). Women appear to be more risk-averse, which validates previous findings (Eckel and Grossman 2008; Ioannou and Sadeh 2016). The expected discount rate decreases with age ($\lambda_{\delta-age} = -0.065$; $s.e. = 0.028$). This contradicts previous research showing that older individuals require higher discount rates as they tend to be more concerned about receiving the payoffs within their lifetime (Chesson and Viscusi 2000). Our results could be due to the limited age range in our sample. The minimum age is 20 and maximum age is 37, with a mean at 23.89.

5.3.2 Discrete heterogeneity—Cluster analysis

To explore the extent of heterogeneity in risk attitudes and intertemporal choices in the sample, we use K-means clustering and group the decision makers in our sample based on the individual level CRRA estimates and discount rates. After eliminating the outliers from this analysis (two subjects with discount rates higher than 50%), we identify three types of decision makers, as shown in Fig. 3. Even though the scree plot does not give a clear picture with respect to the number of clusters to analyze

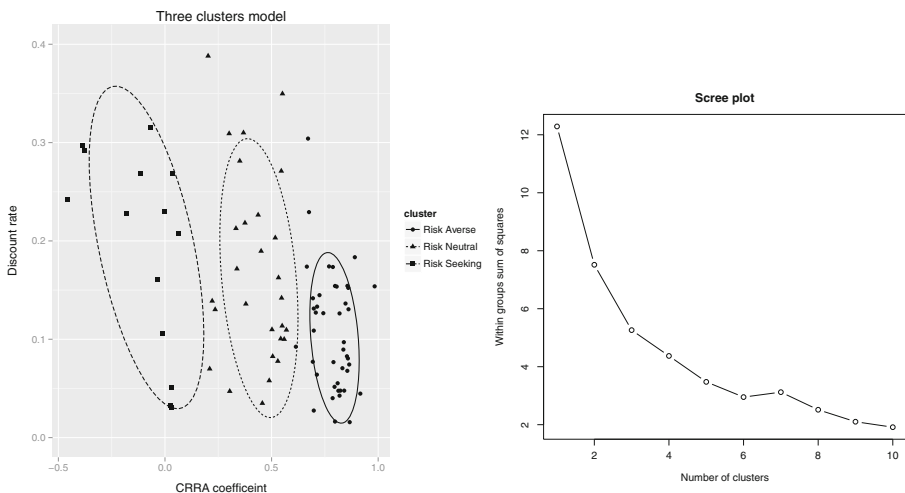


Fig. 3 Three-cluster model and scree plot

(either 3, 4, or 5 types are acceptable), the ease of interpretation for three groups outweighs the gains in variance explained (7%) from the 4-cluster model over the 3-cluster model. Table 6 reports the descriptive statistics of the three groups.

Group 1 (n=14) consists of risk-seeking participants (mean CRRA is -0.10) who are impatient and require a high discount rate (19.49% on average) to postpone consumption. We call this group the “low patience” type. The second group (n=28) includes moderately risk-averse subjects (mean CRRA is 0.42), who require a mean interest rate of 16.94% to postpone consumption. This group represents individuals with average patterns of behavior on both risk and time dimensions. We label this group the “moderate patience” type. The third group (n=43) embodies very risk-averse (mean CRRA is 0.79) individuals who require a low discount rate to postpone consumption (10.75% on average). We name this group the “high patience” type.

5.3.3 Cumulative discount rate distributions

To further explore heterogeneity, we investigate the cumulative discount rate distribution for each level of risk decision makers are willing to accept. The cumulative discount rate distribution refers to the percentage of participants who are willing to accept a discount rate lower than or equal to a specific value. Given that they belong to clusters 1, 2, and 3, respectively, we know that decision makers are risk-seeking, risk neutral, and risk-averse and use this information to produce cumulative discount rate distributions for each cluster of participants.

From Fig. 4 and Table 7, we can determine that about 28% of the risk-seeking decision makers accept a discount rate less than or equal to 15%. Sixty-four percent of the risk neutral individuals will accept a discount rate of 20% or less, while only 35% of the risk-seeking individuals will accept such a discount rate. Forty-six percent of the risk-seeking people will only accept a discount rate of 25% or more, i.e. they would reject discount rates below 25%. However, 74% of the risk neutral and 64% of the risk-averse decision makers will accept at least a 20% discount rate, therefore they accept offers between 20% and 25%, otherwise rejected by risk-seeking participants, showing that they would be more patient. Risk-seeking individuals will reject plans or policy measures that are attractive to decision makers belonging to a different type. In this context, identifying decision makers’ risk preferences and adapting to their needs becomes important to ensure a wide acceptance of policy measures involving intertemporal choices.

Table 6 Cluster analysis

	Groups		
	<i>Low Patience</i>	<i>Moderate Patience</i>	<i>High Patience</i>
CRRA	-0.1	0.42	0.79
Discount Rate	19.49%	16.94%	10.75%
Group Size (%)	16%	35%	49%




Fig. 4 Cumulative distribution of discount rates for different types

5.4 Assessing the correlation between risk and time preferences

The hierarchical Bayes specification allows for a direct measure of the correlation between risk and time preferences, which is not possible in maximum likelihood methods. Table 8 reports the covariance matrix V_{θ} that characterizes the unexplained variability in risk and time preferences across subjects. The diagonal elements, showing posterior variances of our model parameters, are higher and show significant heterogeneity between subjects for both risk and intertemporal choices. Therefore,

Table 7 Cumulative distribution of discount rates for different types

	Risk-seeking	Risk-neutral	Risk-averse
accept 15% or less	28.57%	53.57%	74.41%
accept 20% or less	35.71%	64.28%	95.34%
accept 25% or less	64.28%	78.57%	97.67%

there is scope for analyzing such heterogeneity via individual demographic and psychographic variables. The off-diagonal elements indicate similar patterns of behavior across subjects. Although the values are based on the exponentially transformed variables and the size of the covariance is not readily interpretable, the fact that the covariance is significant and negative supports the finding that more risk-averse individuals exhibit increasing patience. The propensity for errors is lower for more risk-averse individuals, since the covariance between the risk aversion parameter r and the noise parameters ν and μ is significant and negative. The propensity for error is higher for more impatient individuals ($\text{cov}(\nu, \delta)=0.425$).

We use the MCMC draws to compute the parameters’ posterior means for each subject in our sample, based on the exponential discounting model. The correlation of these predicted values is -0.42 ($p < 0.01$). Therefore, we find evidence of a negative correlation between risk aversion and impatience.

We also estimate the discount rates under the assumption that subjects are risk neutral, and separately estimate individual level risk preferences. The correlation between discount rates under risk neutrality and CRRA coefficients is 0.16, not significant ($p > 0.05$). This correlation is similar in sign and size with the correlation between the model free discount rates and risk attitudes.

To assess whether the pattern of correlation differs between risk-seeking, risk neutral and risk-averse individuals, we conduct three regression analyses, each using a dataset corresponding to the types reported in Section 5.2.3. We find no significant correlation between risk and time preferences for risk-averse and risk neutral individuals (Pearson’s $r = -0.219$, $p > 0.05$, and Pearson’s $r = -0.217$, $p > 0.05$), and a significant negative correlation between risk and time preferences for risk-seeking decision makers (Pearson’s $r = -0.545$, $p = 0.044$). This shows that the model is able to recover various patterns of correlation in the data. Moreover, it is crucial to correct for the concavity of the utility function before assessing the size and sign of this correlation. We use the risk preferences task to identify the curvature of the utility

Table 8 Covariance matrix of model parameters

	Risk Aversion r	Noise Parameter μ	Discount Rate δ	Noise Parameter ν
Risk	0.204 ^{1a}	-0.159^a	-0.112^a	-0.242^a
Aversion r	(0.032)	(0.047)	(0.044)	(0.063)
Noise		0.560 ^a	0.268 ^a	0.386 ^a
Parameter μ		(0.117)	(0.084)	(0.120)
Discount			0.703 ^a	0.425 ^a
Rate δ			(0.121)	(0.124)
Noise				1.171 ^a
Parameter ν				(0.206)

¹Statistical signif. codes:(a) 99% interval does not include 0

function. At the individual level, we obtain lower discount rates under the assumption of risk aversion than for risk neutrality. However, this individual-level relationship does not impact the correlation between risk and time preferences across subjects (Andersen et al. 2008, page 612).⁵

5.5 Sensitivity to background consumption

Our estimation technique requires an estimate of the background consumption ω to control for the curvature of the utility function, as discussed in Section 4.1. We exogenously set this parameter for the above estimation at $\omega=12$, for comparison with Andersen et al.'s (2008) study. To assess the sensitivity of our results to different levels of background consumption, we vary the value of ω (see Fig. 5). Andersen et al. (2008) find that their estimated discount rates are not very sensitive to the level of background consumption, while the CRRA coefficient increases from 0.67 when $\omega=50$ to 0.82 when $\omega=200$. Andreoni and Sprenger (2012a) report that estimated discount rates double when daily consumption is varied from \$3.52 to \$14.09. Although their methodology to elicit curvature-controlled discount rates is invariant to the level of daily background consumption, when replicating the Andersen et al. (2008) joint estimation technique, Laury et al. (2012) find that discount rates increase from a low of 14.3% when $\omega=0$ to about 20% as ω increases. We find that our estimates of risk and time preferences show little sensitivity to background consumption; the CRRA coefficient increases from 0.49 when $\omega=1$ to 0.57 when $\omega=200$. The discount rates increase from a low of 10.3% for $\omega=1$ to 15.6% for $\omega=200$, similar to the pattern reported in Laury et al. (2012). Hence, there is a slight increase in the discount rate when we increase background consumption.

6 Discussion and conclusions

The elicitation of risk and time preferences is pivotal in understanding individuals' decisions such as health-related choices (e.g., dieting, addictions), as well as in developing efficient public policies that impact private decisions, involving tradeoffs between now and the future (e.g. environmental policies, incentives for education). In this paper, we present a method to simultaneously elicit and estimate decision makers' risk attitudes and intertemporal choices. We jointly elicit risk and time preferences, as doing so provides significantly lower discount rates than in experiments which assume subjects to be risk neutral, and closer to what is a priori considered to be reasonable discount rates. We use a novel individual level estimation procedure based on a hierarchical Bayes methodology, which embeds standard discounting functions and risk specifications from the decision making literature. The

⁵To check whether the negative correlation is driven by our choice of model, we conducted a simulation study to gauge whether we can assess different patterns of correlation. We simulate data from the model, imposing positive, negative, or no correlation between risk and time preferences. We recover all patterns of correlation (simulated correlation = {0.38, -0.03, -0.38} vs. estimated correlation = {0.40, 0.09, -0.21}, respectively). We thank an anonymous reviewer for raising this point.

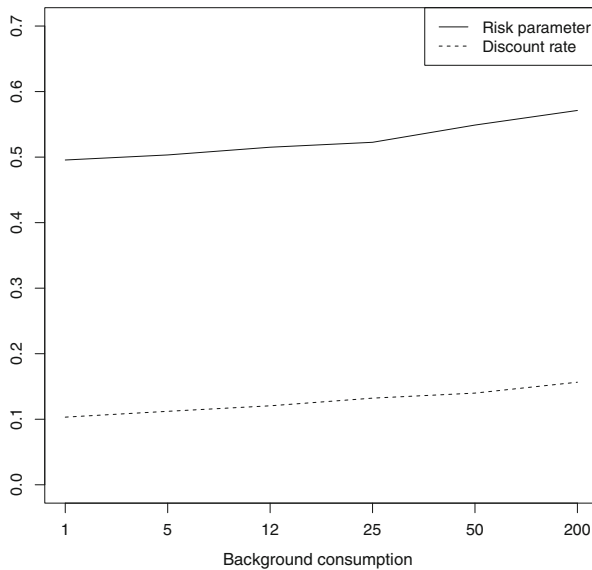


Fig. 5 Estimated discount rates and CRRA coefficients as a function of daily background consumption

method accounts for similarities among individuals and provides robust estimates of the model parameters without ignoring the interdependency between subjects or over-weighting their individual differences. We achieve this by pulling individual estimates towards a group mean. This effect is more pronounced when estimates are less reliable, mitigating the noise in the estimation. Hierarchical Bayes modeling is particularly useful when the number of subjects in an experiment is relatively large, while the data from each subject are few or noisy, which is the case in many decision making experiments including ours. The methodology allows us to obtain more accurate estimates, without increasing the duration of the experiment or the number of tasks.

Using our hierarchical Bayes methodology to jointly estimate risk and time preferences, we recover the probability distributions for the group-level CRRA coefficients and discount rates, as well as individual level risk attitudes and intertemporal choices. We find that, on average, individuals require an annual discount rate of 12% to postpone consumption and are moderately risk-averse (the constant relative risk aversion coefficient is 0.515). The results are in line with previous literature on risk and time preferences and further support the need to control for the curvature of the utility function when estimating discount rates in experiments. We also test the robustness of our results against alternative specifications for time preferences, e.g. Mazur’s hyperbolic and Prelec’s hyperbolic discounting models, as well as a mixture specification. The models are easily integrated within our joint estimation of risk and time preferences, another advantage of the hierarchical Bayes methodology. While Prelec’s hyperbolic specification fits the data better, the qualitative results are similar to those implied by exponential discounting. This implies that correcting for the curvature of the utility function is important in estimating discount rates.

We use the individual level parameter estimates to assess the level of heterogeneity in risk and time preferences. We find that women tend to be more risk-averse, and people tend to be more patient with age. We conduct a cluster analysis and find three types of decision makers that vary in terms of their risk attitudes and intertemporal choices. Risk-averse individuals are more patient, while risk-seeking individuals tend to be more impatient. This is a consequence of the reported negative correlation between risk attitudes and intertemporal choices. The results of the model also suggest that there is significant unobserved heterogeneity, worth assessing in future studies. Policy makers can use data on preference elicitation to offer type-specific plans or design mechanisms that produce desirable self-selection, for public policies involving health issues (e.g., smoking, dieting), education requirements, or environmental plans. An interesting avenue for further research would be to investigate the most suitable ways to present information on public policy measures to different types of decision makers.

Reassuring for environmental policy makers, Ioannou and Sadeh (2016) find that discounting is similar for monetary and environmental decisions. Using an experimental design where subjects make choices between monetary rewards vs. planting bee-friendly plants, the authors show that while discount rates are not statistically significantly different, people tend to be more risk-averse for environmental decisions. Ioannou and Sadeh's (2016) study sequentially elicits risk and time preferences and reports no correlation between subjects' discount rates and risk aversion, a result we also confirm when eliciting risk and time preferences sequentially. Their study differs from ours in a few important respects. We correct for the concavity of the utility function when eliciting discount rates, and use a hierarchical Bayes methodology to *simultaneously* estimate risk and time preferences. Also, Ioannou and Sadeh's (2016) study involves real stakes, while we use hypothetical stakes. More research is necessary to assess the correlation between risk and time preferences across different domains, elicitation protocols and estimation methods.

Acknowledgments We would like to thank seminar participants at the Rotterdam School of Management, Erasmus University and the University of Texas at Dallas for helpful comments and suggestions. We are particularly grateful to Stefano Puntoni, Jason Roos, W. Kip Viscusi, Peter Wakker, and an anonymous referee for their insightful comments.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Abdellaoui, M., Bleichrodt, H., l'Haridon, O., & Paraschiv, C. (2013). Is there one unifying concept of utility? An experimental comparison of utility under risk and utility over time. *Management Science*, 59(9), 2153–2169.
- Anderhub, V., Güth, W., Gneezy, U., & Sonsino, D. (2001). On the interaction of risk and time preferences: An experimental study. *German Economic Review*, 2(3), 239–253.

- Andersen, S., Harrison, G.W., Lau, M.I., & Rutström, E.E. (2008). Eliciting risk and time preferences. *Econometrica*, 76(3), 583–618.
- Anderson, L., & Stafford, S. (2009). Individual decision-making experiments with risk and intertemporal choice. *Journal of Risk and Uncertainty*, 38(1), 51–72.
- Andreoni, J., & Sprenger, C. (2012a). Estimating time preferences from convex budgets. *American Economic Review*, 102(7), 3333–3356.
- Andreoni, J., & Sprenger, C. (2012b). Risk preferences are not time preferences. *American Economic Review*, 102(7), 3357–3376.
- Camerer, C. (2004). Individual decision making. In Kagel, J.H., & Roth, A.E. (Eds.), *Handbook of Experimental Economics*. Princeton University Press (pp. 587–616).
- Cardenas, J.C., & Carpenter, J. (2013). Risk attitudes and economic well-being in Latin America. *Journal of Development Economics*, 103(C), 52–61.
- Chesson, H.W., & Viscusi, W.K. (2000). The heterogeneity of time-risk tradeoffs. *Journal of Behavioral Decision Making*, 13, 251–258.
- Chesson, H.W., & Viscusi, W.K. (2003). Commonalities in time and ambiguity aversion for long-term risks. *Theory and Decision*, 54, 57–71.
- Coller, M., & Williams, M. (1999). Eliciting individual discount rates. *Experimental Economics*, 2(2), 107–127.
- Conte, A., Hey, J.D., & Moffatt, P.G. (2011). Mixture models of choice under risk. *Journal of Econometrics*, 162, 79–88.
- Eckel, C.C., & Grossman, P.J. (2008). Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior & Organization*, 68(1), 1–17.
- Falk, A., Meier, S., & Zehnder, C. (2013). Do lab experiments misrepresent social preferences? The case of self-selected student samples. *Journal of the European Economic Association*, 11(4), 839–852.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351–401.
- von Gaudecker, H.M., van Soest, A., & Wengstrom, E. (2011). Heterogeneity in risky choice behavior in a broad population. *American Economic Review*, 101(2), 664–94.
- Holt, C.A., & Laury, S.K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644–1655.
- Ioannou, C.A., & Sadeh, J. (2016). Time preferences and risk aversion: Tests on domain differences. *Journal of Risk and Uncertainty*, 53(1), in press.
- Jackman, S. (2009). *Bayesian Analysis for the Social Sciences*. Wiley Series in Probability and Statistics. Wiley.
- Jarnebrant, P., Toubia, O., & Johnson, E. (2009). The silver lining effect: Formal analysis and experiments. *Management Science*, 55(11), 1832–1841.
- Kang, M.J., Rangel, A., Camus, M., & Camerer, C.F. (2011). Hypothetical and real choice differentially activate common valuation areas. *Journal of Neuroscience*, 31(2), 461–468.
- Kass, R.E., & Raftery, A.E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795.
- Kimball, M.S., Sahm, C.R., & Shapiro, M.D. (2009). Risk preferences in the psid: Individual imputations and family covariation. Working Paper 14754, National Bureau of Economic Research.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112, 443–477.
- Laury, S.K., McInnes, M.M., & Swarthout, J.T. (2012). Avoiding the curves: Direct elicitation of time preferences. *Journal of Risk and Uncertainty*, 44(3), 181–217.
- Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics*, 107(2), 573–597.
- Mazur, J.E. (1987). An adjustment procedure for studying delayed reinforcement. In Commons, M.L., Mazur, J.E., Nevin, J.A., & Rachlin, H. (Eds.), *The Effect of Delay and Intervening Events on Reinforcement Value*. Erlbaum (pp. 55–76).
- Meyer, A. (2013). Estimating discount factors for public and private goods and testing competing discounting hypotheses. *Journal of Risk and Uncertainty*, 46(2), 133–173.
- Nilsson, H., Rieskamp, J., & Wagenmakers, E.J. (2011). Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, 55(1), 84–93.

- Prelec, D. (2004). Decreasing impatience: A criterion for non-stationary time preference and hyperbolic discounting. *Scandinavian Journal of Economics*, 106(3), 511–532.
- Rossi, P., Allenby, G., & McCulloch, R. (2005). Bayesian Statistics and Marketing. No. 13 in Wiley series in probability and statistics, Wiley.
- Rossi, P.E., & Allenby, G.M. (2003). Bayesian statistics and marketing. *Marketing Science*, 22(3), 304–328.
- Taylor, M. (2013). Bias and brains: Risk aversion and cognitive ability across real and hypothetical settings. *Journal of Risk and Uncertainty*, 46(3), 299–320.
- Toubia, O., Johnson, E., Evgeniou, T., & Delquié, P. (2013). Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science*, 59(3), 613–640.