Internal and External Validity of Experimental Risk and Time Preferences

Christian Belzil Modibo Sidibé

École Polytechnique/
CREST/ CYRANO/ IZA

Duke University/
CREST

October 31, 2016

ERID Working Paper Number 237

Economic Research Initiatives at Duke
WORKING PAPERS SERIES

Eruditio et Religio

Internal and External Validity of Experimental Risk and Time Preferences*

Christian Belzil[†] Modibo Sidibé [‡]

October 31, 2016

Abstract

Using a unique field experiment from Canada, we estimate individual preference over risk and time and show considerable heterogeneity in both dimensions and relatively stable distributions across our various specifications, which include hyperbolic, quasi-hyperbolic discounting as well as subjective failure probability over future payments. We investigate the prediction power (transportability) of the estimated preference parameters when used to explain the take-up decision of higher education grants where financial stakes are approximately seven to fifty times larger than the cash transfers used to elicit preferences. We find that both long-run discount factors and subjective payment failure risk parameters have a high degree of transportability across tasks, while parameters characterizing short-run discount preferences are irrelevant when considering higher-stakes decisions.

Keywords: discounting, risk aversion, time inconsistency, transportability.

JEL codes: C91,D12 ,D81

*We thank Attila Ambrus, Eric French and Victoria Prowse, as well as seminar participants at many conferences and seminars for their comments and discussions. Nicolas-Aldebrando Benelli provided outstanding research assistance.

[†]CREST, CNRS, Ecole polytechnique, Université Paris-Saclay, IZA and CIRANO; christian.belzil@polytechnique.edu

[‡]Duke University and CREST; modibo.sidibe@duke.edu

Introduction

Characterizing individual choices when confronted with the temporal allocation of resources under uncertain states of nature is essential to answering policy questions. As such, preferences for risk and time are at the center of most important economic decisions. For instance, the saving decision of workers, the education choice of teenagers, and the decision to purchase insurance are all examples of economic choices that depend crucially on preferences over time and risk. Individual differences in decisions reflect not only the heterogeneous effects of institutional constraints, but also differences in attitudes toward risk and time. While parameters characterizing preferences for risk and time are theoretically well-defined, they are known to be difficult to identify using observational data. This is particularly true about the rate of time preference which is clearly under-identified within dynamic discrete choices (Magnac and Thesmar, 2002). Not surprisingly, this has led many economists to use various experimental designs to estimate discount factors and risk-aversion parameters.

Despite the substantial experimental literature on risk and time preference, very little is known about the robustness of these parameters under different non-testable modeling assumptions and their validity in different contexts. We use the specificity of a Canadian artefactual (field) experiment carried out in several high schools between 2008 and 2009 to infer the empirical distribution of deep-preference parameters and evaluate their capacity to explain individual choices in a different context. The elicitation of risk and time preference is based on three distinct designs. Since the experiment incorporates various combinations of inter-temporal choices with and without front-end delays, we examine a classical exponential discounting framework, as well as more-general settings allowing for time-inconsistent preferences. These include a popular time-inconsistent specification of individual preferences the $\beta - \delta$ (Laibson, 1997), the quasi hyperbolic formulation due to Loewenstein and Prelec (1992).[1]. Finally, we estimate a subjective risk model where individuals perceive a potential difference in payment realization between immediate and future payments.

As the experiment incorporates one segment designed to elicit deep preference parameters and a second one designed to infer the value attached to higher education financial aid opportunities, we can perform two separate tasks. We first estimate the distribution of risk and time preference parameters from in-

[1] In the rest of this paper, we refer to these as the quasi-hyperbolic and hyperbolic discounting models following Cohen et al. (2016)

dividual tasks designed to elicit them (as is usually done in the experimental literature) and investigate their internal validity. Second, we evaluate the external validity of those parameters. To do so, we use the second segment of the experiment, in which young individuals choose between immediate cash payments and higher education financial aid packages that are worth between 7 to 50 times the cash payments of the first segment (those used to estimate preference parameters).

The capacity to transport scientific results obtained from an experimental setting to different environments is one of the most controversial topics in data-driven sciences. The issue is particularly acute in the experimental economics literature since preferences for risk and time are often elicited in laboratories using a set of tasks that entail relatively low stakes and also rely on homogeneous sub-populations. As a consequence, the transportability of elicited preferences to other contexts is a contentious issue.[2]

We find that under classical exponential discounting, there is a relatively low level of relative risk aversion among young individuals, and a relatively high level of annual discount factors. Estimates obtained from values of background consumption indicate that an increase in background consumption shifts both the distributions of risk aversion and discount factors to the left illustrating the direct and indirect impact of background consumption on the curvature of the utility function.

We quantify the relative importance of pure heterogeneity and shocks in shaping these distributions, and show that relative dispersion of preference parameter to random noise almost always exceeds 1 for all questions highlighting the limitations of the pure expected utility theory. Allowing for time-inconsistent behaviour, we find that 25% of high school students are subject to present-bias in the quasi-hyperbolic model, while more than 73% of students in the sample are found to be time-inconsistent, according to the hyperbolic discounting model.[3] Allowing for subjective payment probability, we show that individuals assign a relatively low (high) probability of payment (non-payment) when choosing a later cash payment.

The hyperbolic discounting outperforms the quasi-hyperbolic specification,

[2] A similar but narrower concept exists in the experimental literature, and is referred to as the magnitude effect (Andersen et al., 2013). Our notion of transportability encompasses potential contextual validity, while the magnitude effect is more concerned by the stability of parameters across tasks of different rewards. The issue is analyzed in (Dohmen et al., 2011).

[3] This contrasts with the results from Andersen et al. (2014) who do not find any evidence of hyperbolic discounting.

as well as the subjective risk model, in all criteria used for internal validity purposes suggesting that the simplicity of the $\beta - \delta$ formulation comes at the expense of flexibility.

Finally, when we consider the transportability the estimates of risk and time preferences, we find that subjective-risk interpretation of present-biased behavior appears to be more credible than those obtained by assuming time-inconsistent preferences. The subjective payment failure probabilities estimated from low-stakes decisions are found to be the most important determinant of higher-stakes decisions regarding educational financial aid; and the quasi-hyperbolic model approach to low-stakes decisions leads to a distribution of present-bias parameters that appear to be disconnected from higher education financing decisions.

In section 1, we briefly review the literature on estimating preferences for risk and time and the links between the econometric literature on external validity and the notion of transportability that has recently emerged in the experimental literature. In Section 2, we provide a detailed description of the field experiment In Section 3, we present the behavioral model used when modeling the first phase of the experiment. Section 4 is devoted to our estimation results. We discuss the notion of transportability of both preference and subjective-risk parameters in Section 5.

1 Literature

1.1 Estimating Preferences for Risk and Time

The literature on estimating preferences for risk and time is vast and heterogeneous. The inherent difficulties in identifying deep structural preference parameters using observational data have led many economists to use various experimental designs to estimate discount factors and risk-aversion parameters.

The initial experimental literature on estimating discount factors has relied heavily on the "Multiple List (MPL)" approach, with monetary payments and assumed risk neutrality
(Coller and Williams, 1999). In the MPL approach, individuals face a sequence of binary choices between immediate and future payments characterized by increasing interest rates. The point at which individuals revert to later payments provides interval identification of the discount factor. Because this approach

has often led to unrealistic estimates of discount factors, economists have gradually recognized the need for estimating discount factors and risk-aversion parameters jointly. This is exemplified by Andersen et al. (2008), who have designed an experiment that using MPL to capture both preferences over time and risk. They show how assuming linear preferences may bias discount factors (rates) downward (upward).

More recently, Andreoni and Sprenger (2012) designed an experimental method based on "Convex Time Budget," which allows point estimation of preference parameters.

Our approach to estimating preferences, therefore, differs in many respects. Our data incorporate several MPLs for risk aversion (Holt and Laury 2002 and Eckel and Grossman 2008) and time preferences Coller and Williams (1999). In order to use all the information available, we estimate preference parameters as fixed effects and decompose dispersion in individual choices between its noisy component and its true (structural) component. We need neither to eliminate individuals who would revert their choices more than once nor to discard any of the lists. Our approach also allows us to estimate the degree of cross-sectional dispersion in background consumption and to examine the sensitivity of the parameters of the distribution of risk-aversion estimates (location and scale) to various assumptions.

1.2 Model Validity and Contextual Transportability

Another objective of our study is to evaluate the contextual transportability of our measures of preferences. This is a highly contentious issue. In the recent statistical literature, Bareinboim et al. (2012) have developed a theoretical analysis of the necessary conditions for transportability. They use statistical notions popular in the artificial intelligence literature to translate the conditions of transportability into a formal analytical framework. In econometrics, the term "transportability" is rarely analyzed, but related issues are often discussed in conjunction with the notion of external validity. Most researchers concerned with external validity have used various social experiments in order to evaluate the forecasting performance of a model estimated on a given population (a control group) and have used its parameters to predict the behavior of a different population.

For instance, Rosen (1985) analyzed the predicted impacts of a housing subsidy program using estimates obtained externally from a population that had no access to housing subsidies. Wolpin and Todd (2006) followed a similar approach. Using data from a large-scale government program in Mexico, they estimated a dynamic behavioral model of parental decisions about children's schooling and fertility among households not covered by the program and used their estimates to forecast behavior of households affected by it.

While economists and social scientists are aware of the potential drawbacks of the experimental approach (pre-test effects, post-test effects, Rosenthal effects, framing, strategic manipulation), the extent to which specific results may be transposed between contexts is rarely investigated at a formal level. One particular area of economics in which the issue of transportability is acute is the experimental economics literature, as preferences for risk and time are often elicited within laboratory experiments. Because the vast majority of studies infer those parameters using a set of tasks that entail relatively low stakes on homogeneous sub-populations (for instance, college undergraduates), the transportability of elicited preferences to other contexts has raised serious doubts.

One specific area where the notion of transportability (taken generally) has raised interest is in Game Theory. In experiments devoted to the measurement of Level-K strategic behavior, the stability (or persistence) of the distribution of levels across different games has recently attracted some attention. For instance, experimental economists debate the existence of individual-specific traits that may translate into a form of "strategic sophistication" [4].

2 The Field Experiment

The experiment was conducted from October 2008 to March 2009 on a sample of 1,248 Canadian full-time students aged from 16 to 18 years, drawn from both urban and rural sites across Canada. Supplementary information on the experiment can be found in Belzil et al. (2016) or Johnson and Montmarquette (2015).

All subjects were presented with the full set of decisions and were paid for one, randomly selected, at the end of the session. The subjects were informed that they would be paid for one decision, but they did not know which one at the beginning of the session. The questions were split into three groups. First, the subjects answered a set of questions aimed at measuring their rate of time preference. The second set of questions related to the measurement of risk atti-

[4] Recent examples include Camerer et al. (2004) and Georganas et al. (2015).

tudes. The third group of questions consisted of a sequence of choices between a cash payment, to be paid within one week of the day of the experiment, and a specific financial aid package covering educational expenses.

In this section, we describe the sets of questions used to infer the distribution of preference parameters. The last set of questions will be described in the section devoted to transportability.

2.1 Time Preferences

The first part of the experiment, which consists of 48 questions, is designed to identify individual time preferences. The experiment is based on multiple price lists (MPLs) with monetary payments (Coller and Williams, 1999). The interest rate increases monotonically in a price list, such that the point at which individuals switch from preferring earlier payments to later payments carries interval information about their intertemporal preferences. When individuals revert from the earlier to the later payment at most once, this approach induces a narrow-bracketing of the discounting range.

In addition, the experimental protocol manipulates not only the number of time periods between the earlier and the later payment, but also the timing of the earlier payment. This sort of "front-end delay" is a well known method to elicit time-inconsistent preferences but also implies access to multiple pieces of information (measurements) on an individual-specific preference parameter. This provides strong arguments for using a factor representation of the true preference parameter.

In Table 1, we report the different payments along with their associated timing. All decisions imply a choice between an earlier payment, denoted c, and a later payment, denoted d , to be paid $t + \tau$ weeks from the day of the experiment. Individuals in the experiment were presented payoff tables such as the one illustrated in Table 1, with six symmetric intervals. For example, the first decision involves a choice between a payment of $75 to be paid within one day and a future option, $75.31 to be paid in one month.

Figure 1, which uses information on the first reversion to evaluate discount rates, describes an extensive variation in the relative interest rates. For example, choices between payments received in one day vs. one month, imply that 30% of the sample has a monthly interest rate over 200%, while only 13% has a yearly interest rate over 200%. Similarly, interest rates measured at four months indicate that less than 5% of the sample has an interest rate between 10-20%,

Table 1: Discount Rates Lotteries

Panel 1: 48 Choices

# questions	Horizon 1	Horizon 2
6	1 day	1 month
6	1 day	1 year
6	1 week	1 week + 1 month
6	1 week	1 week + 1 year
6	1 month	2 months
6	1 month	1 year + 1 month
6	3 months	4 months
6	3 months	1 year + 3 months

Panel 2: Payment options and interest rate

Payment ($) 1 day	Payment ($) 1 month	Payment ($) 1 year	Annual Interest(%)
75	75.31	78.75	5
75	75.63	82.5	10
75	76.25	90.00	20
75	78.13	112.5	50
75	81.25	150.0	100
75	87.5	225.0	200

Notes: (i) There are eight sets of choices. (ii) Each choice set is composed of six questions implying different individual interest rates.

while choices at one year and three months suggest that more 15% is within that range. Although the implied discount rates are lumpy, the experiment measures discount rates at several points on the discounting curve. In addition to information on each MPL, the coherence of the total set of MPLs provides identification content. As a consequence, MPL with delays in both instantaneous and distant payments are crucial in matching the discounting behavior of individuals and, therefore, identifying the parameters that control the discounting functions.

Figure 1: Implied interest rates

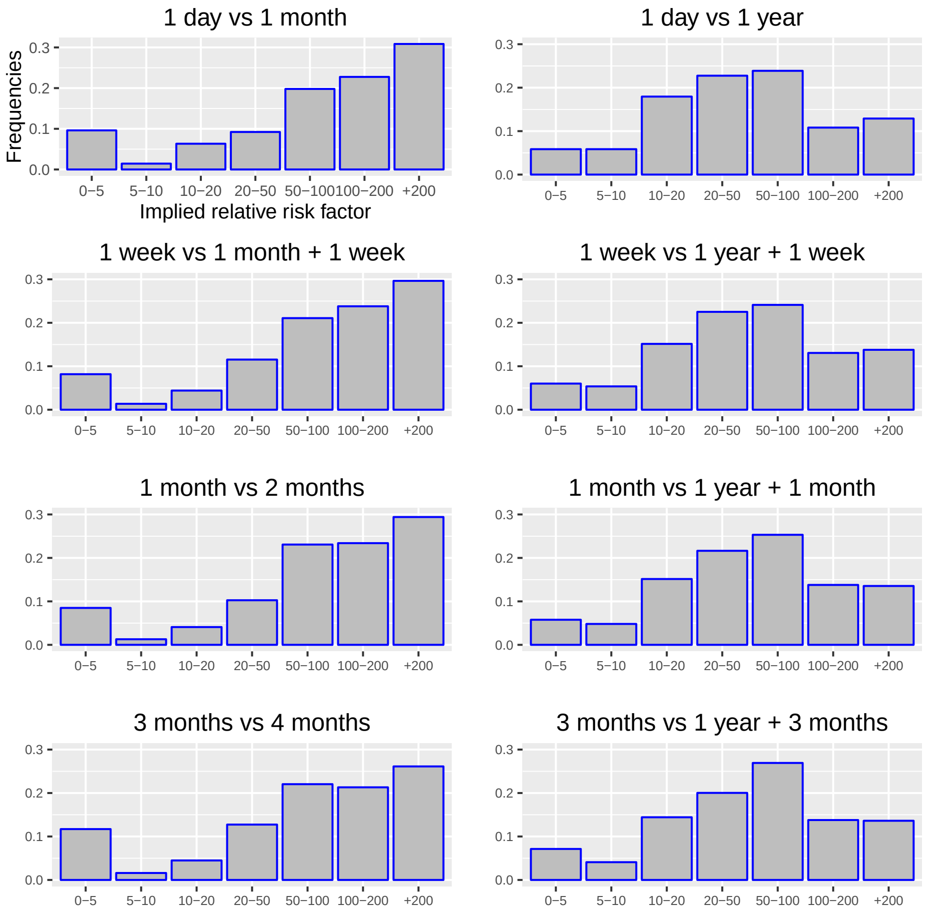

2.2 Risk Aversion

The measurements of risk aversion are based on two distinct sets of experimental procedures. The first uses the Holt and Laury (2002) MPL approach with decreasing objectively stated risks. The second one uses the mechanism popularized by Eckel and Grossman (2008). Both strategies consist of choosing between a lottery with a given payoff distribution and one characterized by more extreme payoffs. In both cases, the cutoff point at which an agent switches from the "average" to the extreme lottery, is identified. The major difference between the two strategies lies in the fact that while the first one pins down a cut-off probability, the second one identifies a cut-off payoff.

Table 2 describes the first set of lotteries, which contains three distinct set of lotteries, consisting of 10 choices each. Each choice is binary, and each option is characterized by a low payoff (L), a high payoff (H) and a probability distribution over L and H. For all three different sets, the expected outcome of the first option is higher for the first four lotteries and lower for the last six riskier choices. For instance, this implies that a risk-neutral agent should choose the first option when probabilities of the high payoff are lower than 0.5 and then switch over to choose the second alternative when probabilities are higher than or equal to 0.5.

Table 2: Risk-aversion lotteries

			Choice 1		Choice 2		Choice 3	
			L= 32	L=2	L=24	L=1.5	L=40	L=2.5
	Probabilities		H=40	H=77	H=30	H=57.75	H=50	H=96.25
	L	H	$\mathbb{E}V$	$\mathbb{E}V$	$\mathbb{E}V$	$\mathbb{E}V$	$\mathbb{E}V$	$\mathbb{E}V$
1	0.90	0.10	32.80	9.50	24.60	7.12	41.00	11.87
2	0.80	0.20	33.60	17.00	25.20	12.75	42.00	21.25
3	0.70	0.30	34.40	24.50	25.80	18.38	43.00	30.63
4	0.60	0.40	35.20	32.00	26.40	24.00	44.00	40.00
5	0.50	0.50	36.00	39.50	27.00	29.62	45.00	49.38
6	0.40	0.60	36.80	47.00	27.60	35.25	46.00	58.75
7	0.30	0.70	37.60	54.50	28.20	40.88	47.00	68.12
8	0.20	0.80	38.40	62.00	28.80	46.50	48.00	77.50
9	0.10	0.90	39.20	69.50	29.40	52.12	49.00	86.88
10	0.00	1.00	40.00	77.00	30.00	57.75	50.00	96.25

Notes: (i) $\mathbb{E}V$ for expected value, L for Low payoff, H for High payoff. ii) Payoffs are in Canadian \$. iii) Source: SRDC-CIRANO Field Experiment on Education Financing.

Similarly to choices regarding discount rates elicitation, it is possible to define bounds on risk aversion for an individual endowed with a CRRA utility function. Figure 2 reports the implied bounds of risk-aversion factors for each set of lotteries. We find that there is substantial variation in the risk-aversion rates both across individuals and within each set of ten lotteries. For instance, around 11% of individuals have a relative risk-aversion rate between -0.14 and 0.15 when using the first set of lotteries; when using the 3rd set of lotteries, the proportion of individuals with a risk aversion rate between -0.14 and 0.15 increases to 17%. Although these differences may appear small in magnitude, they clearly show that MPL are not redundant, and illustrate clearly the sources of identification.

Figure 2: Turning points and implied relative risk -aversion factors

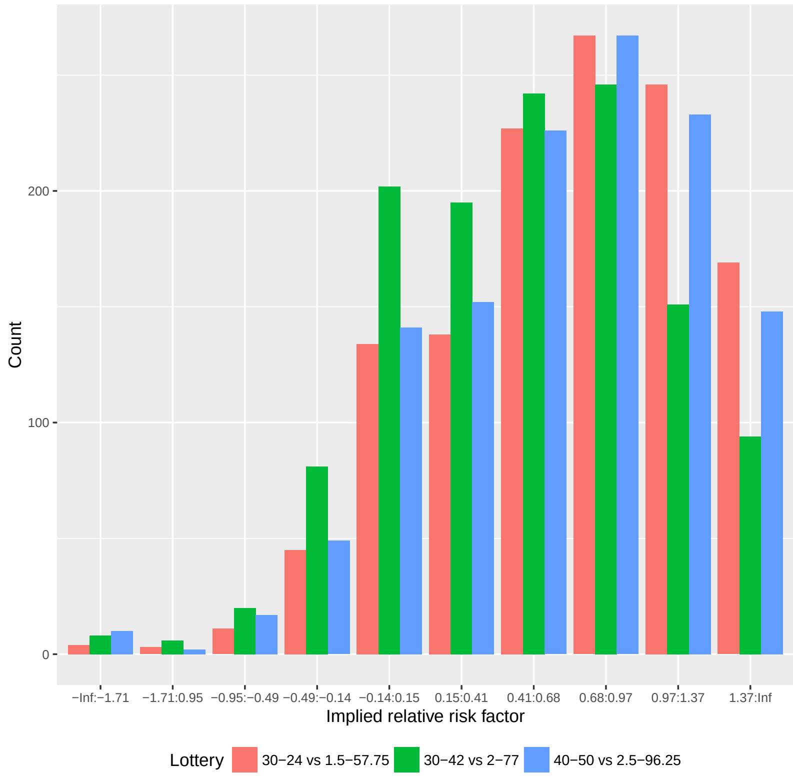

The second set of risk-aversion lotteries uses the gamble mechanism of Eckel and Grossman (2008). Subjects are presented with five gambles, each entailing five binary choices. In all cases, individuals choose between two options. In each case, the first option has lower dispersion than the second and also has a lower expected payoff. Each option has a 50% probability of being drawn. As is well known, this method produces extremely reliable estimates of risk preference (Charness et al., 2013).

The first three gambles include a sure payoff of $48, while the last one introduces variation in the amount of the sure alternative—$42 and $54 respectively. The structure of the lotteries is as follows: from an initial state where the first option offers a sure payoff, alternatives within a specific choice move in different directions such that the dispersion of payoffs increases. As a consequence, the last decision of each lottery involves choices with extreme payoffs. For example, the third lottery includes a choice between $(8, 104)$ and $(0, 112)$. The implied risk-aversion distribution is reported in Figure 3. In contrast to the results obtained using the approach of Holt and Laury (2002) and already disclosed in Figure 2, this approach implies different risk-aversion rates for each lottery.

Table 3: Gamble lotteries

Lottery 1				Lottery 2				Lottery 3			
Choice 1		Choice 2		Choice 1		Choice 2		Choice 1		Choice 2	
48	48	40	64	48	48	42	66	48	48	38	62
40	64	32	80	42	66	36	84	38	62	28	76
32	80	24	96	36	84	30	102	28	76	18	90
24	96	16	112	30	102	24	120	18	90	8	104
16	112	8	120	24	120	16	128	8	104	0	112

Lottery 4				Lottery 5			
Choice 1		Choice 2		Choice 1		Choice 2	
42	42	36	60	54	54	44	68
36	60	30	78	44	68	34	82
30	78	24	96	34	82	24	96
24	96	18	114	24	96	14	110
18	114	10	122	14	110	6	118

3 The Model

We present the framework used to estimate preferences for risk and time. For generality, we describe a specification that allows for time-inconsistent preferences using the $\{\beta, \delta\}$ framework (Laibson, 1997).

At time 0, when the experiment takes place, the intertemporal utility of consumption stream $\{c_0, c_1, ...c_\tau..\}$ is equal to

$$V_0 = U(c_0) + \beta \sum_{t=1}^{T} \delta^t U(c_t),$$

where $U(.)$ is the per-period utility function; β is the present-bias parameter; and δ measures the classical discount factor. Throughout this paper, we assume that the per-period utility, denoted $U(.)$, belongs to the Constant-Relative-Risk

Figure 3: Gamble implied relative risk-aversion factors

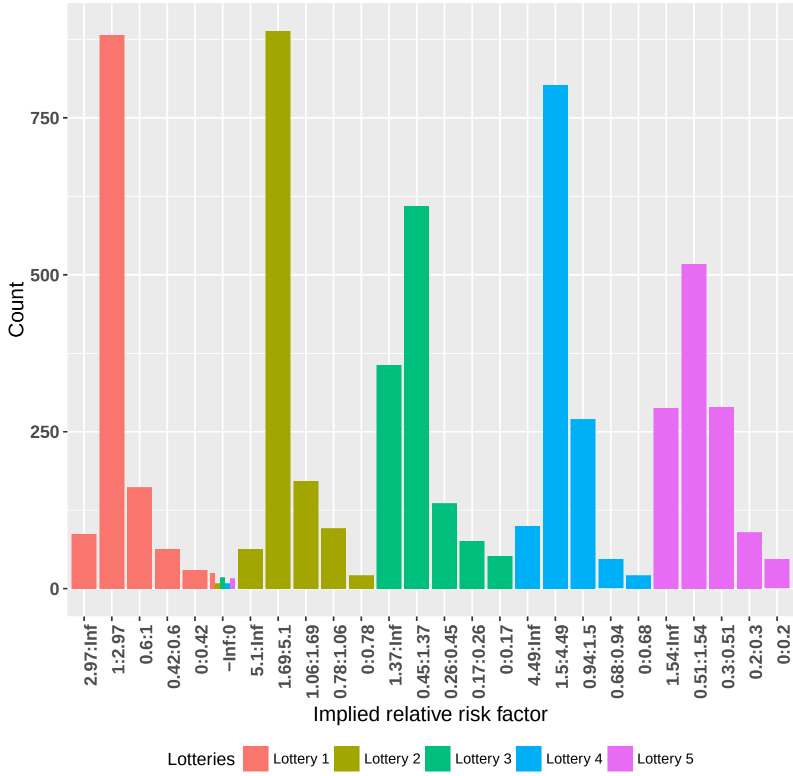

Aversion (CRRA) family; that is:

$$U(c_t;\theta) = \frac{c_t^{1-\theta}}{1-\theta}\,.$$

Under the assumption that individuals do not smooth consumption when offered relatively small amounts, the windfall cash payment is completely consumed in that period.

Because most payments are indexed in months, we do the following adjustments: any payment to be made m months from the time of the experiment, the relevant individual discount factor, $\delta(m)$, is expressed as:

$$\delta(m) = \frac{1}{1+\frac{m}{12}\cdot r},$$

where r is the individual discount rate of interest. Any payment offered within a week is treated as immediate. Finally, the classical exponential discounting version of the model is obtained after setting β to 1 for all individuals.

To be realistic, we must take into account that each individual is endowed with a background (reference) level of consumption denoted w that represents the optimal level of consumption in the absence of any windfall cash payment. We are agnostic about the origins of w - it may depend on parents' income or part-time employment participation.

3.1 Measuring Risk Preferences

In this part of the experiment, individuals exercise (55) binary choices. Each decision, indexed by r, requires a choice between two lotteries. In each case, the second lottery is unambiguously more risky than the first one. The first lottery is characterized by a low payoff, c_{lr}, and a high payoff, denoted c_{hr}, while the second lottery entails a high payoff denoted d_{hr} and a low payoff denoted d_{lr}. For each lottery, the probability of the high outcome is equal to p_{hr}, and the probability of the low outcome is p_{lr}.

The utilities of the less risky lottery and the more risky, denoted $V_1(r)$ and $V_2(r)$, respectively, are equal to

$$\begin{aligned} V_1(r) &= p_{hr} \cdot U(w + c_{hr}) + (1 - p_{hr}) \cdot U(w + c_{lr}) \\ V_2(r) &= p_{hr} \cdot U(w + d_{hr}) + (1 - p_{hr}) \cdot U(w + d_{lr}). \end{aligned}$$

The probability of choosing the second lottery is given by

$$\Pr\{\eta_r > V_2(r) - V_1(r)\} = 1 - \Phi(V_2(r) - V_1(r)), \tag{1}$$

where η_r is an idiosyncratic error term that is independent and identifically distributed across individuals and questions. This $\Phi(.)$ is the cdf of a normal distribution $\mathcal{N}(0, 1)$.

3.2 Measuring Time Preferences

In order to identify and estimate time preferences, we use the 48 choices between an early and a later cash payment. As is common in the literature on eliciting time preferences (Andersen et al., 2008), we assume that individuals do not smooth consumption when offered relatively small amounts of money, such as those offered in the first phase, and, therefore, consume the windfall cash payment in that period. This implies that binary choices may be modeled using only the properly discounted utilities associated with the relevant periods.

Put generally, the per-period utility in period 0 of choosing any amount a to be paid at time t is equal to

$$\begin{cases} U(w+a), & \text{if } t = 0 \\ \beta\delta^t U(w+a), & \text{if } t = 1, 2, \ldots, T. \end{cases}$$

Each potential choice q consists of two mutually exclusive cash payments—c_q at time t and d_q at time $t+\tau$. In the absence of consumption smoothing induced by the cash payments, choosing between c_q at time t and d_q at time $t+\tau$ boils down to comparing two distinct sums of utilities. To introduce noisy measurements of preference parameters, we assume an additive idiosyncratic error, which is i.i.d. across questions and across individuals, and denoted by ε_{iq}.

Setting the indicator y_q to 1 when option 1 (c_q) is chosen and to 0 when option 2 (d_q) is chosen, the expression for the probabilities of choosing the earlier payment are given by the following expressions:

$$\Pr\{y_q = 1\} = \begin{cases} \Pr\left\{\varepsilon_q > (U(w) - U(w+c_q) + \beta\delta^\tau[U(w+d_q) - U(w)\right\}, & \text{no front-end delay} \\ \Pr\left\{\varepsilon_q > \beta\delta^t[U(w) - U(w+c_q] + \beta\delta^{t+\tau}[U(w+d_q) - U(w)]\right\}, & \text{front-end delay,} \end{cases} \tag{2}$$

3.3 Subjective Payment Risk

In many circumstances, the incidence of present-bias behavior may be confounded with the asymmetric effect of risk on the utility of accepting a cash payment in the present as opposed to the future. This issue has been raised in conjunction with Prospect theory, by Halevy (2008), and has been investigated empirically in Andreoni and Sprenger (2012).

In the current experiment, we cannot rule out the possibility that some individuals perceived a positive probability that payments offered in the future may never be realized and, therefore, decided to put more weight on current cash transfers. At the same time, the existence of a subjective risk of non-payment is not likely to affect choices between payments both taking place in the future.

To investigate this issue, we revert to a standard exponential discounting model and cast the model in a framework in which nearly-immediate cash payments (to be received within one week) are treated as certain outcomes, but future ones are perceived as risky. Formally, we assume the following:

- *Payments offered within one week:* All individuals assume that cash payments offered within one week are going to be realized with probability equal to 1.

- *Payments offered beyond one week*: All individuals assume that payments offered beyond one week from the time of the decision will be realized with probability p and will not be paid with probability $1-p$.

- *Heterogeneity in subjective risk:* As we did for the model with time-inconsistent preferences, we assume that each individual is endowed with his/her specific subjective payment probability p.

- *Stationary Subjective Payment Probabilities:* The individual-specific payment probabilities are independent of the duration between the current period and the theoretical payment period.

Given these assumptions, the expression for the probability of choosing the earlier payment in absence of front-end delay becomes

$$\Pr\left\{ \varepsilon_q > \begin{cases} p \cdot \delta^{\tau} U(w+d_q) + (1-p) \cdot \delta^{\tau} U(w) - U(w+c_q) - \delta^{\tau} U(w)\}, & \text{no front-end delay} \\ \delta^{t} U(w) + p \cdot \delta^{t+\tau} U(w+d_q) + (1-p) \cdot \delta^{t+\tau} U(w) & \\ \qquad -p \cdot \delta^{t} U(w+c_q) - (1-p) \cdot \delta^{t} U(w) - \delta^{t+\tau} U(w)\}, & \text{front-end delay} \end{cases} \right\} \tag{3}$$

3.4 Estimation

We estimate the model by maximum likelihood techniques. The likelihood function is the product of all individual choice probabilities. Each individual choice history is itself composed of the product of 48 probabilities related to preferences for times, which is multiplied by the product of 55 probabilities that relate to choices between lotteries. The likelihood function of the discounting and the risk-aversion decisions, which are denoted $\mathcal{L}^d$ and $\mathcal{L}_i^r$, respectively, are equal to

$$\begin{aligned} \mathcal{L}_i^d &= \prod_{q=1}^{48} \Pr\{y_{iq}=1\}^{y_{iq}} \cdot \Pr\{y_{iq}=0\}^{(1-y_{iq})} \\ \mathcal{L}_i^r &= \prod_{r=1}^{55} \Pr\{y_{ir}=1\}^{y_{ir}} \cdot \Pr\{y_{ir}=0\}^{(1-y_{ir})} \end{aligned}$$

Denoting $\Theta_i = \{\beta_i, \delta_i, \theta_i, p_i\}$ as the set of parameters to be the estimated vector to be estimated for each individual i, the total likelihood, denoted $\mathcal{L}$, is the product of all individual contributions and is simply equal to

$$\mathcal{L}(\Theta) = \prod_{i=1}^{1248} \mathcal{L}_i^d \cdot \mathcal{L}_i^r$$

The model is optimized using a standard numerical optimization algorithm. Since there is no cross-sectional relationship between individuals, we estimate the model for each agent.[5]

The identification of risk and time preferences in our setting comes from the Multiple Price Lists. Very high number of questions by individual, parametric assumptions on the utility function, along with the existence of an error term, allow us to point identify an individual-specific risk and discount rate. The main identification challenge will emerge as we consider the estimation of individual background consumption along with risk and time preferences. To illustrate the identification challenge, we consider the decision of an agent when offered a choice c_1 now and c_2 at time τ. Agent chooses c_1 over c_2 when

$$\frac{(w+c_1)}{1-\theta}^{1-\theta} + \delta^\tau \frac{w}{1-\theta}^{1-\theta} > \frac{w}{1-\theta}^{1-\theta} + \delta^\tau \frac{(w+c_2)}{1-\theta}^{1-\theta} \tag{4}$$

Upon examining equation 4, it is clear that the non-linearities involved in the utility components of the decision imply that individual differences in background consumption cannot be eliminated from the the choice probabilities. This observation provides an intuitive explanation for the parametric identifiability of background consumption.

3.5 Internal Validity

An important aspect of our approach is to evaluate the validity of our elicited parameters under various non-testable modeling assumptions. As such, we rely on internal validity, using statistical methods to gauge the relative performance of the various models. We use three (3) criteria to assess the performance of a model. The first criterion is based on the Likelihoods $\mathcal{L}^{(\mathcal{D})}$, $\mathcal{L}^{(\mathcal{R})}$ and $\mathcal{L}$. Obviously, the larger the log-likelihood, the better the model. The second uses

[5]Our initial Monte-Carlo studies indicate that optimizing a model with **1248** parameters is 100 times more burdensome than estimating 1,248 models with one parameter.

information criteria—namely, those of Akaike and Schwartz.[6]

Where k is the number of parameters, and n the number of observations.

The underlying idea is that given finite data, an excessive number of parameters leads to instability. As such, information criteria introduce a penalty for models with an excessive number of parameters. Finally, we use a measure of fit, defined as the sum of squared differences between the empirical and predicted probabilities.

$$\mathcal{Y} = \sum_J (\bar{Y}_j - \widehat{Y}_j)^2 \tag{7}$$

$\mathcal{Y}$ is an analog of sum-mean-square error of the estimator, and as a consequence, it aggregates individual prediction errors into a single measure. In our analysis, we will distinguish between the fit of the discount questions, $\mathcal{Y}^{\mathcal{D}}$, and that of risk-aversion questions $\mathcal{Y}^{\mathcal{R}}$.

4 Results

In this section, we report the main findings of our model. In the first part, we maintain the assumption that individuals use classical exponential discounting. In the second section, we allow for time-inconsistent behavior. Then, we introduce the model with uncertain payment realization. Finally, we consider the transportability of our parameters.

4.1 Exponential discounting

In this subsection, we examine two particular issues: the impact of alternative assumptions about background consumption and the importance of allowing for a noisy measurement of preference parameters.

4.1.1 Sensitivity with Respect to Background Consumption

In the literature, background consumption is rarely estimated. The standard technique consists of calibrating background consumption to an homogenous

[6]We denote them by AIC and BIC, and they are defined by

$$\begin{aligned} \mathcal{I}_{\text{Akaike}} &= 2(k - \log(\mathcal{L}(\zeta))) & (5) \\ \mathcal{I}_{\text{Schwartz}} &= -2\log(\mathcal{L}(\zeta)) + k\log(n), & (6) \end{aligned}$$

sample-wide level, as in Andersen et al. (2008), or estimates a single parameter, as in Andreoni and Sprenger (2012). To illustrate the importance of the assumption on background consumption, we start our analysis by setting the background consumption level to various amounts (\$5, \$10, \$15, \$20 and \$100) in order to evaluate the sensitivity of risk aversion and discount factors. Table 4 reports the resulting distributions of annual discount factors and relative risk aversion, respectively.

First, we note that risk aversion is relatively low, as the median level ranges between 0.58 with the \$5 background consumption level and 0.51 with the \$100 level. The median level of risk aversion of young teenagers is, therefore, much lower than what would be observed for logarithmic preferences. The quasi-equality between the mean and the median levels of risk aversion points toward a relatively symmetric distribution. Not surprisingly, the degree of relative risk aversion changes with background consumption, as the results indicate that the median and the mean relative risk aversion parameter decreases uniformly with background consumption. This is explained by the fact that a change in background (reference) consumption also changes the marginal utility of any cash payment, and that the estimated degree of relative risk aversion is bound to vary with it. However, the difference between relative risk aversion obtained with $w = \$5$ and $w = \$100$ (0.51 vs. 0.58) reveals a relatively low level of sensitivity.

A common result reported in the experimental literature is that discount factors tend to be low when preferences are linear but are found to be higher when preferences are concave (Andersen et al., 2008). So, if alternative assumptions about background consumption impact the estimated risk aversion through the overall concavity of the utility function, they are also bound to affect the distribution of discount rates. An interesting question is whether discount factors increase with background consumption, as we already noted that risk aversion decreases with background consumption.

Our results indicate that the distribution of discount factors shifts to the left as background consumption increases. For instance, the median (mean) discount factor with a \$5 background consumption is 0.76 (0.68) to be compared to 0.65 (0.60) with a background consumption of \$100.

There are two important points to stress. First, estimates of discount factors are more sensitive than risk-aversion estimates to the level of background consumption. This is indicated by the 0.11 differential (0.76-0.65) in median discount factors when moving from $w = \$5$ to $w = \$100$. Such a difference would

Table 4: Initial condition effect

	CRRA				
	$w = 5$	$w = 10$	$w = 15$	$w = 20$	$w = 100$
Panel A: Annual discount rate					
Min	0	0	0	0	0
1st De.	0.0154	0.0149	0.0143	0.0138	0.00812
1st Qu.	0.627	0.603	0.587	0.575	0.49
40th Cent.	0.716	0.697	0.68	0.667	0.591
Median	0.759	0.74	0.728	0.718	0.652
Mean	0.684	0.672	0.663	0.654	0.608
Sd	0.279	0.278	0.278	0.278	0.281
60th Cent.	0.801	0.787	0.777	0.771	0.722
3rd Qu.	0.867	0.857	0.85	0.844	0.807
9th De.	0.94	0.936	0.933	0.931	0.921
Max	1	1	1	1	1
Panel B: Risk-aversion factor					
	CRRA				
	$w = 5$	$w = 10$	$w = 15$	$w = 20$	$w = 100$
Min	-0.571	-0.649	-0.795	-0.866	-0.886
1st De.	0.346	0.354	0.358	0.358	0.33
1st Qu.	0.469	0.471	0.469	0.466	0.417
40th Cent.	0.54	0.539	0.536	0.528	0.477
Median	0.583	0.582	0.577	0.573	0.512
Mean	0.582	0.574	0.561	0.556	0.498
Sd	0.293	0.246	0.163	0.159	0.137
60th Cent.	0.626	0.621	0.613	0.606	0.542
3rd Qu.	0.686	0.677	0.668	0.661	0.588
9th De.	0.777	0.76	0.75	0.741	0.659
Max	5	5	1.2	1.12	0.866
Panel C: Comparison					
	CRRA				
	$w = 5$	$w = 10$	$w = 15$	$w = 20$	$w = 100$
$\mathcal{Y}^{\mathcal{D}}$	1.18	1.2	1.21	1.22	1.27
$\mathcal{Y}^{\mathcal{R}}$	0.749	1	1.22	1.4	2.5
$\mathcal{Y}$	1.93	2.2	2.43	2.62	3.77
$\mathcal{L}^{\mathcal{D}}$	-21226	-21362	-21452	-21527	-21909
$\mathcal{L}^{\mathcal{R}}$	-30291	-31217	-31984	-32607	-36242
$\mathcal{L}$	-51517	-52579	-53436	-54134	-58151
$\mathcal{I}_{\text{Akaike}}$	108026	110150	111864	113261	121293
$\mathcal{I}_{\text{Schwartz}}$	132397	134521	136235	137632	145664

imply sizable effects in any behavioral model.

Second, the impact of background consumption on discount factors is intricate. Despite its negative impact on risk aversion, increasing the level of background consumption reduces discount factors. This illustrates the impact of background consumption on the curvature of the utility function both directly and indirectly through its effect on the estimated parameter of relative risk aversion.

4.1.2 Noise in Preference Measurements

In line with standard methods in factor analysis, we quantify the importance of stochastic shocks. To do this, we decompose total utility between structural parameters (factors) and noise. For any given choice s the difference in expected utility between each relevant option, V_s, has the form $V_s = V_2(s) - V_1(s)$.

For the discount rate questions, $V(s)$ is defined as

$$V_s = U(w) + \beta\delta^{\tau}U(w + d_s) - U(w + c_s) - \beta\delta^{\tau}U(w),$$

while for risk-aversion questions, $V(s)$ is equal to

$$\begin{aligned} V_s \quad = \quad & p_{hs} \cdot U(w + d_{hs}) + (1 - p_{hs}) \cdot U(w + d_{ls}) \\ & - p_{hs} \cdot U(w + c_{hs}) + (1 - p_{hs}) \cdot U(w + c_{ls}). \end{aligned}$$

It is simple to obtain a factor-noise ratio for each choice s by comparing the standard deviation of $V(s)$ to the standard deviation of the noise component, which is set to 1.

$$\frac{\sigma(V(s))}{\sigma(\varepsilon_s)}.$$

The difference in expected utility for each choice s depends on individual preference parameters and on the monetary value of lotteries. In total, we computed 103 factor-noise ratios. To ease presentation, we focus on choices that disclose balanced proportions and, therefore, avoid extreme choices (those for which one specific option tends to be chosen by almost everyone). In practice, this implies selecting choices that are located in the middle of various lists.

In Table 5, we report the results obtained for a subset of eight discounting questions. The first panel is devoted to discounting questions, and the second to risk-aversion questions. We find that the discount rate factor-noise ratio fluctuates more across questions than the risk-aversion pendant. For discount factors, the ratio ranges between ten and one, but for most questions, differences in discount factors (and risk aversion) appear more important than random shocks. For risk aversion, the range is much smaller, as the ratio is between four and one. Overall, and for those questions reported, we find that differences in risk aversion are about two times as important as noise.

While preference heterogeneity is driving a majority of choices, we take these ratios as evidence in favor of the factor model specification. For instance, when considering the choice between $75 in 30 days and $78.13 in 60 days, our estimates imply that 50% of total utilities are explained by purely random noise. Similarly, for two of the risk-aversion questions selected—(34,82) vs. (24,96) and (28,76) vs. (18,90)—random noise appears almost as important as differences in risk aversion.

Table 5: Variance Decomposition

Panel A: Annual discount rate		
Choices in $	Horizon in days	Std. Dev V(s)
75 vs 78.13	1 vs. 31	9.5
75 vs 78.13	7 vs 37	3.4
75 vs 78.13	30 vs. 60	1.17
75 vs 78.13	90 vs 120	0.92
75 vs. 90	1 vs 361	11.1
75 vs. 90	7 vs 367	6.5
75 vs. 90	30 vs 390	5.3
75 vs. 90	90 vs 450	4.6
Panel B: Risk-aversion factor		
Choices in $	Probabilities	Std. Dev V(s)
(40,32) vs (77,2)	0.5, 0.5	3.1
(30,24) vs (57.75,1.5)	0.5, 0.5	2.1
(50,40) vs (96.25, 2.5)	0.5, 0.5	4.3
(32,80) vs. (24,96)	0.5, 0.5	2.0
(36,84) vs (30,102)	0.5, 0.5	2.6
(28,76) vs (18,90)	0.5, 0.5	1.5
(30,78) vs. (24,96)	0.5, 0.5	2.6
(34,82) vs. (24,96)	0.5, 0.5	1.4

4.1.3 Estimating the Distribution of Background Consumption

Because our estimation method is based on all questions available in the experiment, one natural extension is to attempt to estimate the distribution of background consumption. As we already showed evidence that background consumption affects the estimated level of risk aversion, one can also suppose that ignoring dispersion in background consumption may affect the degree of cross-sectional heterogeneity in risk aversion and discount factors. For this reason, we re-estimated the model, adding background consumption as an extra individual-specific quantity. We impose a minimum level of consumption to avoid numerical problems. The distribution of discount factors, risk aversion and background consumption is reported in the first three columns of Table 6.[7]

First, our estimates of background consumption indicate that a vast majority of young students (more than 60%) use a background consumption reference point that approaches 0. At the same time, the distribution is skewed to the right, as the average level is around \$19. Second, and as expected, the estimates indicate that the estimation of background consumption has affected the distribution of discount factors more than it has affected risk aversion. To see this, it is sufficient to compare the distribution of preferences reported in Table 4 (obtained at various levels of background consumption) with those in Table 6. Both the mean and median level of risk aversion are equal to 0.56, while the median discount factor has increased to 0.80. This latter result, indicating that more than half of the high school student population is endowed with a discount rate below 20% per year, appears especially consistent with standard assumptions in calibrated macroeconomic models.

4.1.4 An Alternative Approach

The relatively low level of background consumption raises obvious numerical problems. As the CRRA may sometimes be defined only for strictly positive-valued consumption levels (depending on the curvature parameter), a very low background consumption level translates into estimating a parameter that lies near the boundary of the parameter space and may, therefore, involve numerical problems.

This observation leads us to define an alternative approach. As noted earlier, the standard assumption is that individuals consume their endowment in a single period, but this still requires individuals to base their decisions not only on the difference between the discounted utilities of the later and current cash payment, but also on the difference between the utilities of consuming their reference level in the future vs. immediately. Put more formally in terms of equations 1 and 2, the take-up probability $\Pr\{y_q = 1\}$ depends not only on $\beta\delta^{\tau}U(w + d_q) - U(w + c_q)$, but also on $U(w) - \beta\delta^{\tau}U(w)$, which can not be evaluated when w=0.

[7] Estimates around the minimum are reported as being approximately equal to 0.

Table 6: The Distribution of Discount Factors, Risk Aversion and Background Consumption

Panel A: Standard Model and Alternative Model

	Standard model			Alternative model	
	Discount	Risk aversion	Back. Cons.	Discount	Risk Aversion
1st decile	0.190	0.325	$\simeq 0$	0.086	0.313
1st quart.	0.656	0.429	$\simeq 0$	0.685	0.450
40th dec.	0.761	0.510	$\simeq 0$	0.781	0.529
Median	0.802	0.561	$\simeq 0$	0.816	0.582
Mean	0.716	0.560	19.31	0.718	0.628
St Dev	0.271	0.198	45.15	0.282	0.522
60th Cent.	0.836	0.610	$\simeq 0$	0.846	0.630
3rd Quartile	0.887	0.688	5.37	0.896	0.712
9th Decile	0.955	0.788	120.9	0.942	0.854

Panel B: Comparison

Internal Validity	Standard model	Alternative model
$\mathcal{Y}$	1.67	1.65
$\mathcal{L}$	-47184	-50144
$\mathcal{I}_{\text{Akaike}}$	101855	105280
$\mathcal{I}_{\text{Schwartz}}$	138412	129651

Notes: (i) The Standard Model refers to the specification where background consumption is treated as an individual-specific parameter.

To circumvent this issue, we estimate an alternative model based on the normalization condition that $U(w) \equiv 0$, along with the assumption that $w \approx 0$. For very small background consumption levels, individual decisions are practically based solely on utilities evaluated at the relevant cash transfer level, and the choice probabilities are simply:

$$\begin{aligned} \Pr\{y_q = 1\} &= \Pr\{\varepsilon_q > \beta\delta^\tau U(\tilde{c} + d_q) - U(\tilde{c} + c_q)\} \\ &\approx \Pr\{\varepsilon_q > \beta\delta^\tau U(d_q) - U(c_q)\}, \end{aligned} \tag{8}$$

Estimating this equation eliminates the problem of dealing with background consumption parameters at the boundaries of the parameter space and may, therefore, be regarded as an interesting alternative to the standard model. In what follows, we re-estimate our model using equation 8 and refer to it in Table 6 as the alternative model.

Columns 4 and 5 of Table 6, report a similar distribution of discount factors across the standard and alternative models. For instance, the median discount factor obtained with the alternative model is equal to 0.82, while the standard approach leads to a median of 0.80. We obtain similar estimates at practically every quantile.

For risk aversion, the main difference between the distributions obtained under the the standard and alternative models lies at the right tail of the distribution. While both medians are comparable (0.56 vs. 0.58), the alternative model implies a mean relative risk aversion (0.63) substantially higher than the one obtained under the standard approach (0.56) and, therefore, reveals a distribution of risk aversion skewed to the right.

One interesting question is whether the alternative model can fit the data as well as the standard model. Because the alternative model has fewer parameters, we use the measures of internal validity described earlier.

In panel B of Table 6, we report a few statistics to evaluate the relative capacity to fit the data. For both likelihood and mean-squared errors-based measures, we report an overall statistic and a decomposition by discount-rate and risk-aversion questions. The results suggest that the alternative model is capable of fitting the data as well as the standard Model. Based on a mean squared errors principle, the standard model performs slightly better, but this is not the case when using likelihood-based principles. While the standard approach leads to a higher likelihood, the larger number of parameters also implies a significant penalty and leads to the domination of the alternative model when using the information criteria.

4.2 Time-inconsistent Preferences

A vast literature in economics and experimental psychology has documented the incidence of individual behavior violating standard exponential discounting.[8] The view that many individuals are subject to "preference reversal" and discount future events with a declining rate of time preference is now widespread in the economics profession. This has led economists to propose a "hyperbolic discounting" specification of the basic intertemporal decision model. The most common parametrization of time-inconsistent preferences is the two-parameter model in which immediacy is captured in a single parameter (usually denoted β) and in which the standard time-preference parameter is represented by a classical discount factor (usually denoted δ).[9]

Because identifying present bias requires specific designs, a large body of empirical evidence has been obtained in laboratory experiments. While the

[8] Frederick et al. (2002); Cohen et al. (2016) review the literature in depth.

[9] See Laibson (1997) for more details.

vast majority of studies are non-structural, the more recent literature has gradually moved toward the estimation of structural preference parameters. This is achieved by Benhabib et al. (2010), as well as Andreoni and Sprenger (2012), who estimated individual-specific present bias and discount factor parameters, along with risk-aversion parameters. Interestingly, both Benhabib et al. (2010) and Andreoni and Sprenger (2012) find practically no evidence of dynamically inconsistent preferences.[10]

We now extend our approach to take into account potential time-inconsistent preferences. As a first step, we use the $\beta - \delta$ specification of Laibson (1997) and estimate three different parameters: the risk aversion factor (θ); the long run discount factor (δ); and the short run discount factor (β). The discounting function, denoted $D_i(t)$, is

$$D_i(t) = \beta_i \cdot \delta_i^t. \tag{9}$$

The literature on time-inconsistent preferences has considered alternative functional forms of the discounting function, notably, the axiomatic derivation of Loewenstein and Prelec (1992), which is based on a generalized hyperbolic function:

$$D_i(t) = (1 + \alpha_i \cdot t)^{-\frac{\gamma_i}{\alpha_i}}, \tag{10}$$

where α_i and γ_i are positive-valued parameters. The discount function converges to the classical continuous time exponential discounting function ($\exp(-\gamma t)$) when $\alpha_i \to 0$ and approaches a step function when $\alpha_i \to \infty$. We refer to this specification as the hyperbolic discounting function.

In all cases, we consider specifications that ignore heterogeneity in background consumption which is set to \$20. For the sake of comparison with the existing literature, we also estimate the quasi-hyperbolic models assuming linear utility. The results are found in Table 7. We first comment on the estimates from the quasi-hyperbolic specification of Laibson (1997). Two results are important to stress. First, the estimates under a CRRA utility function show that about 25% of the population display "present-biased" behavior. Indeed, 10%

[10] Models of time-inconsistent preferences have also been applied to observational data such as consumption and saving (Laibson, 1997), welfare participation (Fang and Silverman, 2009), job search (DellaVigna and Paserman, 2005), health-club contracts (DellaVigna and Malmendier, 2006) but identification remains problematic, as identifying declining discount rates requires at least two parameters. Indeed, most papers rely on reduced-form implications of hyperbolic discounting but avoid estimating the corresponding structural parameters. When estimated structurally, hyperbolic discounting parameters inferred from observational data are usually assumed to be homogeneous, as in Fang and Silverman (2009).

Table 7: The Distribution of Time-inconsistent Preferences

Panel A: Quasi-hyperbolic

Utility Function	CRRA			Linear		
	δ	β	1st year Discount	Disc. Fac.(δ)	P. Bias(β)	1st year Discount
1st decile	0.01	0.73	0	0.01	0.50	0.261
1st quart	0.57	$\approx$1	0.0003	0.51	0.66	0.476
40th dec.	0.67	$\approx$1	0.021	0.95	0.82	0.504
Median	0.71	$\approx$1	0.128	$\approx$1	0.86	0.647
Mean	0.65	0.94	0.415	0.76	0.79	0.58
St Dev	0.28	0.19	0.45	0.36	0.23	0.254
60th Cent.	0.77	$\approx$1	0.538	$\approx$1	$\approx$1	0.665
3rd Quartile	0.84	$\approx$1	1	$\approx$1	$\approx$1	0.813
9th Decile	0.93	$\approx$1	1	$\approx$1	$\approx$1	0.887

Panel B: Hyperbolic

Utility Function	CRRA			Linear		
	γ	α	1st year Discount	γ	α	1st year Discount
1st decile	0.11	$\approx$0.00	0.009	0.19	$\approx$0.00	0.015
1st quart	0.46	$\approx$0.00	0.498	0.49	$\approx$0.00	0.0575
40th dec.	1.07	3.48	0.643	0.89	0.023	0.086
Median	1.64	6.72	0.685	1.34	1.40	0.116
Mean	2.32	26.4	0.613	1.98	3.30	0.239
St Dev	2.19	44.0	0.299	1.91	35.0	0.278
60th Cent.	2.23	11.3	0.749	1.96	4.74	0.153
3rd Quartile	3.7	25.5	0.825	2.73	10.1	0.302
9th Decile	5.54	108	0.927	5.30	42.6	0.694

Panel C: Model comparison

Measures of Fit	Q-hyperbolic Discounting	Hyperbolic Discounting	Exponential Discounting
$\mathcal{Y}^{\mathcal{D}}$	1.22	0.767	1.22
$\mathcal{Y}^{\mathcal{R}}$	1.47	1.49	1.4
$\mathcal{Y}$	2.71	2.26	2.62
$\mathcal{L}$ (all questions)	-53737	-51703	-54134
$\mathcal{L}^{\mathcal{D}}$ (discount factor)	-21433	-18860	-21527
$\mathcal{I}_{\text{Akaike}}$	114961	110894	113261
$\mathcal{I}_{\text{Schwartz}}$	151518	147450	137632

Notes: (i) All models have been estimated with a background consumption level fixed at $20.

of the high school students are endowed with a present-bias parameter smaller than 0.73. The median high school student is, therefore, discounting according to the standard exponential model. This obviously implies a relatively high mean value of β (0.94). At the same time, the median annual discount factor, which is equal to 0.71, is relatively close to the median obtained when assuming exponential discounting (with a CRRA and a background consumption set at \$20).

A second important finding, related to the sensitivity of discount factor estimates on the form of the per-period utility function, is the high sensitivity of present-bias behavior to the assumed curvature of the utility function. As Table 7 suggests, the prevalence of present-bias behavior among most young individuals increases significantly when assuming risk neutrality. Specifically, under linear preferences, we find that between 50% and 60% of the population is endowed with a short-run discount factor below 1, with a median parameter around 0.86 to be compared to 1 with a CRRA.

We now turn to the parameter estimates obtained with the hyperbolic discount function. At the outset, it should be clear that, unlike the quasi-hyperbolic specification, it is harder to assign a strict interpretation to the parameter (α and β). For instance, the parameter measuring deviation from exponential discounting (namely, α) has a much wider range (essentially 0 to infinity) than the present-bias parameter of the quasi-hyperbolic specification, which is between 0 and 1, and also affects the ratio of discount factors applicable to future payments. For this reason, we need not expect the same proportion of individuals with a value of β equal to in the quasi-hyperbolic specification as the proportion of individuals with an α close to 0 in the hyperbolic specification.

The estimates obtained with the hyperbolic specification are found in panel B of Table 7. When assuming a CRRA utility function, we find that only 27% of high school students behave according to a standard exponential discounting ($\alpha \approx 0$). This proportion is much smaller than the one obtained with the $\beta - \delta$ model, which was almost 75%. This illustrates how sensitive inference about time-inconsistency is to the functional form of the discount function.

As already noted with the quasi-hyperbolic specification model, the estimated proportion of individuals behaving according to the exponential discounting model is equally strongly affected by the curvature of the utility function. With a linear utility function, about 40% of the population is predicted to be time-consistent, and both the median and average values of α (equal to 1.40 and 3.30, respectively) are substantially lower than their respective values obtained with the CRRA (6.72 and 26.4).

Figure 4: Individual Discounting functions

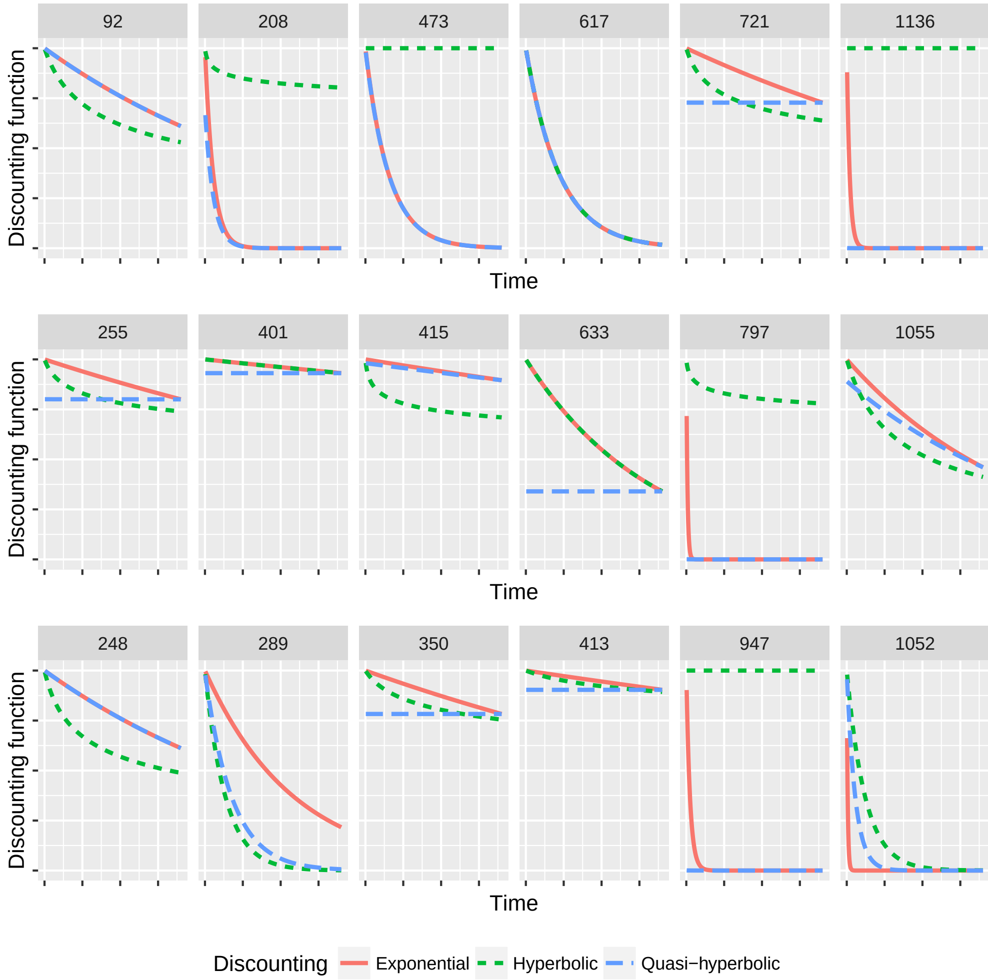

To further illustrate the heterogeneity in individual discount rates, we report in Figure 4 the discounting functions for a random set of 18 individuals. These plots report a diversity of shapes in the discounting functions. For most individuals, the quasi-hyperbolic and hyperbolic discounting functions are not distinguishable. Also, on the sample of individuals, the exponential discounting does well at matching the shapes of discounting functions. When the discounting functions differ, these discrepancies can be small, as with individuals 617, but also extreme 633. For example, individual 1136 has a drastic-step hyperbolic discounting function, while exponential and quasi-hyperbolic discounting functions are essentially flat at around 1.

In Panel C of Table 7, we perform a comparison between quasi-hyperbolic and hyperbolic specifications. For the sake of comparison, we also report the same measures for the classical exponential model described in Table 4.

Regarding measures based on both mean-squared error principles and on likelihood principles (including Akaike and Schwarz criteria), the conclusion is that the hyperbolic specification outperforms the quasi-hyperbolic specification by a wide margin. This is exemplified by the important likelihood differentials between the two specifications when considering either the entire set of questions (2431) or when considering only the discount-factor question (2573).

Interestingly, comparison between the quasi-hyperbolic and the classical exponential model specifications shows that the latter appears to be better. Despite reaching a higher likelihood, the classical exponential discounting performs better on other criterias as the difference in likelihood is not compensated by the number of parameters.

4.3 Time-inconsistent Preferences Vs Subjective Risk

In this section, we evaluate the distribution of subjective payment probabilities that can rationalize observed choices, and we compare it with the distribution of the present-bias parameters. The distributions of relative risk aversion, discount factors and payment probabilities are found in Table 8.

First, we note that individuals assign a relatively low (high) probability of payment (non-payment) when choosing a later cash payment, as both the median and the mean payment probabilities are equal to 0.37. This very high level of subjective risk is easily explained by the relatively low level of relative risk aversion. Logically, any individual endowed with a level of relative risk aversion as low as 0.56 (our estimate of the median) would require a relatively high risk of not being paid in order to refuse future cash payments. In other words, it is tempting to regard the subjective-risk model with mean payment probability equal to 0.37 as the mirror image of the quasi-hyperbolic discount model with a mean present-bias parameter equal to 0.94.

Second, allowing for a probability of non-payment translates into a shift in the overall distribution of discount factors. The median discount factor, now equal to 0.875, is relatively higher than the discount factor obtained in the quasi-hyperbolic discount, model which was equal to 0.71. The potential explanation is related to the fact that when subjective payment risk is accounted for a low discount factor is no longer needed to explain why some individuals systematically prefer immediate (or nearly immediate) cash payments since risk aversion explains part of the desire for immediate rewards.

Third, there is less heterogeneity in the subjective payment probabilities than in discount factors. This is exemplified by the standard deviation of the subjective payment probabilities 0.04 to be compared to 0.34 for discount factors.

Fourth, it is also informative to compare the level of cross-sectional dispersion in payment probabilities to the level of dispersion in short-run discount factor β (its pendant parameter in the quasi-hyperbolic discount model). While β and p are both ranging between 0 and 1, the standard deviation of the short-run discount factor (β) is five times larger than the standard deviation in subjective payment probabilities. Although the quasi-hyperbolic model and subjective-risk representation may be regarded as indistinguishable ways to explain the observed bias for immediate outcomes, the quasi-hyperbolic discount model seems to require substantially more heterogeneity than the subjective-risk model.

We now analyze whether the quasi-hyperbolic formulation and the exponential discounting specification with subjective risk are really similar models. Put differently, we ask whether individuals who regard payment as intrinsically risky within the subjective-risk model are the same individuals endowed with a high level of short-term discount rate in the quasi-hyperbolic model.

We use the distribution of individual-specific parameters and measure their weighting rank correlations across models. Our reasoning is as follows. If the quasi-hyperbolic discounting formulation and the exponential discounting model with subjective risk were identical, we would find a coincidence between the identity of those who are endowed with a low subjective payment probability and a low value of β in the $\beta - \delta$ model. In the panel B of Table 8, we report the rank correlations between the parameters of the quasi-hyperbolic discount model (denoted $\beta^{PB}, \delta^{PB}, \theta^{PB}$) and those obtained from the subjective-risk model ($P^{SR}, \delta^{SR}, \theta^{SR}$).

Our estimates indicate that there is very little correlation between individual payment probabilities (P^{SR}) and present-bias parameters (β^{PB}). The correlation, which is not significantly different from 0, is equal to 0.055. This seems to point toward the existence of a fundamental difference between time-inconsistent and subjective-risk models and to illustrate the fragility of economic reasoning based on a representative agent.

In order to understand the fundamental differences between the two models, it is useful to examine other correlations. One key distinction between the two models lies in the different role played by risk aversion. In the $\beta - \delta$ formulation, individuals endowed with a low value of β automatically tend to accept immediate payments. This would be the case even for risk-neutral individuals. In the model with subjective risk, the appeal for an immediate payment is reinforced by risk aversion. As a result, fitting those two different models on the same data is likely to lead to different distributions of discount factors.

To evaluate the validity of our conjecture, we first examine the correlations between discount factors across models. The correlation between δ^{PB} and δ^{SR}, which is equal to 0.086, suggests the existence of different ranking of discount factors when moving from a quasi-hyperbolic discount formulation to a model with subjective risk.

It is also interesting to note that the distribution of relative risk-aversion parameters, even if identified from static decisions, is affected by modeling assumptions. The correlation is equal to 0.77 and is, therefore, sufficiently far from 1 to conclude that the subjective-risk model also affects the individual ranking with respect to relative risk aversion.

To summarize, while it might be tempting to regard the quasi-hyperbolic discount formulation and the subjective-risk model as substitutes, there are fundamental differences between the two models. In particular, it would be erroneous to assume the existence of a nearly one-to-one correspondence between the preference parameter capturing short-run discounting and the individual subjective payment failure probability.

5 Transportability of Elicited Measures of Preferences and Subjective Risk

In the literature, external validity is based on the capacity of the estimated parameters to be generalized when agents face different decisions. In our setting, we want to determine if our elicited parameters can inform us on the determinants of real-life decisions. Namely, we want to address whether risk and time preferences obtained from moderate-stakes choices are relevant predictors of high-stakes decisions. We refer to this as contextual transportability.

To do this, we use questions that involve the choice between a cash payment to be paid within one week of the experiment, and a specific financial aid package covering educational expenses. As stated earlier, this is the most original dimension of this experiment. Indeed, the segment of the experiment in which young individuals choose financial aid packages over immediate cash

Table 8: The Distribution of Discount factors, Payment Probabilities and Risk Aversion

Panel A: Subjective Risk Model with Exponential Discounting

Utility Function	CRRA		
	Disc. Fac(δ)	payment prob. (p)	Risk Aversion (θ)
1st decile	≈ 0	0.364	0.321
1st quart	0.663	0.371	0.433
40th dec.	0.811	0.373	0.519
Median	0.875	0.374	0.562
Mean	0.737	0.375	0.555
St Dev	0.342	0.037	0.249
60th Cent.	0.942	0.375	0.612
3rd Quartile	≈ 1	0.378	0.674
9th Decile	≈ 1	0.407	0.766

Panel B: Correlations between Structural Parameters Across Models

	β^{PB}	δ^{PB}	θ^{PB}	p^{SR}	δ^{SR}	θ^{SR}
β^{PB}	-	0.140	0.453	-0.055	0.122	0.322
δ^{PB}		-	-0.187	0.604	0.086	-0.165
θ^{PB}			-	-0.182	0.060	0.771
p^{SR}				-	-0.321	-0.168
δ^{SR}					-	-0.020
θ^{SR}						-

Panel C: Comparison

Internal Validity	Quasi-hyperbolic Discounting	Hyperbolic Discounting	Exponential Discounting	Subjective Risk
$\mathcal{Y}^{\mathcal{D}}$	1.22	0.767	1.22	1.04
$\mathcal{Y}^{\mathcal{R}}$	1.47	1.49	1.4	1.46
$\mathcal{Y}$	2.71	2.26	2.62	2.5
$\mathcal{L}$	-53737	-51703	-54134	-50861
$\mathcal{L}^{\mathcal{D}}$	-21433	-18860	-21527	-18457
$\mathcal{I}_{\text{Akaike}}$	114961	110894	113261	109210
$\mathcal{I}_{\text{Schwartz}}$	151518	147450	137632	145767

Note: (i) Parameter estimates with the "PB" superscript are those obtained from the quasi-hyperbolic discount model specification, while those with the "SR" superscript are taken from the model with subjective risk. (ii) Correlations are standard deviation weighted rank correlations.

payment is also the most expensive element of the total cost of this experiment, as the prime objective of those who designed it was to uncover the willingness to pay for higher education financial aid.

As will now become clear, the stakes involved in the decision between cash payments and education financing are substantially higher than the amounts used to elicit preferences for risk and time. The financial aid package may incorporate a single loan, a single grant, or a hybrid package offering a combination of both, and it is to be paid conditional on the student enrolling in a full-time program at any higher education institution in the country (within two years).[11] However, in an effort to focus on higher-stakes choices and because loans need to be repaid, we model only the choice between cash payments and grants.

The monetary value of grants, which ranges from \$500 to \$4,000, is about seven to fifty times the cash transfers used to elicit preferences in the first phases of the experiment. The cash alternative to grants ranges between \$300 and \$700 and is five to ten times higher than those offered in the first phase. In order to picture the purchasing power of these grants, it is important to note that the average tuition fees were \$2,180 for Quebec, \$5,667 for Ontario, \$3,228 for Saskatchewan, and \$5,064 for Manitoba over the period considered. Therefore, a \$2,000 grant would have covered 65% of the total fees at the University of Western Ontario and Queen's University, and almost 100% at McGill University and the University of Montreal.[12] As documented in Belzil et al. (2016), and as normally expected, the grant take-up rates were much higher than the loan take-up rates.

In an ideal setting, transportability would be evaluated by comparing two sets of structural preference parameters or payment probabilities estimated independently. However, this is not possible for many reasons. First, the choice between current cash payments and potential financial aid is obviously not designed to measure present bias since all cash payments are offered at an identical period (one week from the experiment). Second, and regardless of time inconsistencies, because different individuals may have different valuations of financial aid opportunities and different expectations about relevant future outcomes for a wide range of reasons, it is not possible to design an econometric model that would incorporate heterogeneous expectations regarding all rele-

[11] Loans conditions were similar to those of the Federal Canadian Student Loan Program. In monetary terms, cash alternatives varied from \$25 to \$700, while grants and loans varied from \$500 to \$4,000. More details may be found in Johnson and Montmarquette (2015) and Belzil et al. (2016).

[12] These universities are among those that attract most of the elite students in eastern Canada.

vant future outcomes. It is, therefore, very unlikely that we would be able to estimate risk and time preferences as accurately as we could for the first phase of the experiment.

For these reasons, we proceed differently. Instead of estimating a structural model, we estimate a discrete choice model of accepting a higher education grant in place of a current cash payment and allow it to depend on estimated measures of risk and time preferences or subjective payment probabilities from the first step. The foundation of our approach is very simple. If parameters elicited in the earlier phase represent true preferences that enter the decision process for high-stakes choices, they should also play an important role in higher education financing decisions. If not, we should find that they are insignificant.

Before proceeding, two points are important to stress out. First, even if individuals were using the exact same preference parameters as those elicited in the first phase, that would not guarantee a significant effect of risk aversion on the decision to accept a grant. This is easily explained by the fact that individuals who would want to smooth present consumption may also be the ones doing so in the future. As a consequence, one may expect individual decisions to be independent of the curvature of the utility function. On the other hand, individual decisions are most likely affected by individual time preferences.

The second point is related to the possibility that a given fraction of the population of high school graduates may attach no value to higher education financial aid. Ideally, we would like to condition financial aid acceptance on the relevant individual-specific subjective probability of attending higher education. While this is not possible, revealed preferences may be used to identify those who are almost certain not to pursue higher education. In our experiment, around 9% of the students, when faced with a choice between a grant and a current cash payment, never choose the grant. Because one of the grants is as high as $4000, it is reasonable to assume that those individuals will not pursue education beyond high school graduation. For this reason, we exclude the 113 individuals who systematically refuse all grants.

In Table 9, we report financial details and take-up rates of all choices used to investigate the transportability. In total, we estimate seven binary choice models.

To evaluate the transportability of the preference parameters and subjective payment failure probabilities to financial aid questions, we estimate a simple probit model of the probability of accepting a grant as a function of observed

Table 9: Take-up Rates of Various Financial Aid Offers

	Binary Choices		
Choices	**outcome=0**	**outcome=1**	**Take-up rates**
1	Cash: \$25	Grant: \$1000	0.924
2	Cash: \$100	Grant: \$1000	0.909
3	Cash: \$300	Grant: \$1000	0.755
4	Cash: \$700	Grant: \$1000	0.454
5	Cash: \$300	Grant: \$500	0.423
6	Cash: \$300	Grant: \$2000	0.840
7	Cash: \$300	Grant: \$4000	0.919

Notes: Take-up rates measure the fraction of the high school students who accepted the grant in place of the cash payment after excluding all those who systematically refuse all grants.

individual characteristics (family income, education, province,etc.) and pre-estimated preference parameters. We perform this exercise with first-stage estimates obtained under three different models: the classical exponential discounting model (using the $\delta's$ and $\theta's$ reported in Table 6); the quasi-hyperbolic discount model (using the $\beta's$, $\delta's$ and $\theta's$ analyzed in Table 7); and the subjective-risk model. All models are based on the CRRA specification of the per-period utility, and assume a background consumption of \$20. With respect to time inconsistency, we select the quasi-hyperbolic discount model as opposed to the hyperbolic discount model simply because it allows for a clearer distinction between present-bias behavior and classical discounting, although the conclusions drawn from the quasi-hyperbolic discount model are substantially the same as those from the hyperbolic specification.

Defining a binary indicator y_{iq} that is equal to 1 when individual i accepts grant q and 0 if not, the probability that an individual selects the grant is given by:

$$Pr(Y_{iq}=1) = \begin{cases} \Phi(\gamma_q + \gamma_{Xq}\cdot X + \gamma_{\beta q}\cdot \delta_i + \gamma_{\theta q}\cdot \theta_i), & \text{for Exponential Discounting Model} \\ \Phi(\gamma_q + \gamma_{Xq}\cdot X + \gamma_{\beta q}\cdot \beta_i + \gamma_{\delta q}\cdot \delta_i + \gamma_{\theta q}\cdot \theta_i), & \text{for Quasi-hyperbolic Model} \\ \Phi(\gamma_q + \gamma_{Xq}\cdot X + \gamma_{pq}\cdot p_i + \gamma_{\delta q}\cdot \delta_i + \gamma_{\theta q}\cdot \theta_i), & \text{for Subjective Risk Model,} \end{cases}$$

where X denotes a vector of observed characteristics and γ_q is a question-specific intercept term. The marginal effects, along with their estimated standard errors, are reported in Table 10. We now highlight the main findings for each model.

First, there is overwhelming evidence that discount factors obtained when assuming a classical exponential discounting model are transportable to education financing decisions. The positive marginal effects, ranging between 0.47 for the $4,000 grant-$300 cash decision and 0.13 for the $1,000 grant-$300 cash decision, are all significant at 1% and indicate that young individuals who exhibit forward-looking behavior in the presence of small stakes are actually those who are more likely to refuse a cash payment to be paid within one week in order to accept future financial aid. Indeed, discount factors are found to not only be transportable, but also to play a more important role than family financial resources, as the family income marginal effects are lower than the discount factor marginal effects for all seven choices and are, for the most part, insignificant.

The second major finding relates to the effect of time-inconsistent preferences. Specifically, we find no evidence that present-bias behavior displayed in the first phase (with low stakes) is transportable to financial aid decisions. All marginal effects associated with the $\beta_i's$ of the quasi-hyperbolic discount model are insignificant. At the same time, the marginal effects of the long run discount factor (δ) are all positive and significant, and remain quantitatively similar in absolute value to the exponential discounting model, but they are more than ten times as large as the marginal effects of the present-bias parameter.

This finding suggests four possible interpretations. One possibility is that individuals who display present-bias behavior when choices entail low stakes, revert to exponential discounting when faced with larger payoffs.

A second interpretation has to do with the timing of the cash payment, which is announced to take place one week from the time at which decisions are undertaken. If present-bias behavior is caused by immediate temptation only, it is actually possible that offering a cash payment one week in advance is sufficient to remove present-bias behavior among young individuals.

A third interpretation relates to the possibility that some of the young individuals who disclose present-bias behavior in the low-stake decisions use acceptance of financial aid as a form of commitment stimulating higher education participation in the future. In doing so, they would thereby alter their choice set (or at least increase the opportunity cost of some option) when faced with the decision to enter higher education.[13]

A fourth interpretation is simply that seemingly time-inconsistent preferences are instead explained by differences in risk perception. To build on this

[13] A formal presentation of a model where individuals are endowed with preferences ove choice sets may be found in Dekel et al. (2009).

interpretation, we now examine the results from the probit model with the first-stage estimates of the subjective-risk model, which are found in panel C of Table 10. There is overwhelming evidence in favor of the transportability of individual-specific payment probabilities. The marginal effects are large (mostly between 1 and 2) and are all significant at a 1% level. Subjective payment probabilities are, indeed, the dominant factor explaining the decision to accept future loans, implying a less important role of discount factors, although individual discount factors are still significant determinants. In some cases, such as the $1000 grant vs. $100 cash and the $1000 grant vs. $300 cash options, the t-ratios of the discount factor marginal effects are above 1.5 (but below 2.0).

The results may be summarized as follows. When evaluated in terms of their transportability, the estimates of risk and time preferences obtained from a subjective-risk interpretation of present-biased behavior appear to be more credible than those obtained by assuming time-inconsistent preferences. The estimates from the subjective payment failure from low-stakes decisions are found to be the most important determinant of higher-stakes decisions regarding educational financial aid.

On the other hand, the quasi-hyperbolic discounting approach to low-stakes decisions leads to a distribution of short-run discount factors that appears to be disconnected from higher education financing decisions. In this case, only long-run individual discount factors (the δ's) are found to be relevant. Finally, classical discount factors obtained assuming exponential discounting with non-stochastic payment are also found to be highly transportable across stakes.

6 Conclusion

In this paper, we use data from a field experiment to estimate the distribution of deep structural preferences from tasks offering moderate (standard) rewards and investigate their transportability to higher-stakes decisions between immediate cash payments and a higher education fellowship.

We explore the empirical contents of various specifications, including models allowing for present-bias behavior. Our results obtained with a classical exponential discounting model illustrate the sensitivity of risk and time preferences to assumptions regarding background consumption. We show that the hyperbolic discounting of individual preferences due to
Loewenstein and Prelec (1992) provides the best representation of the data (in-

Table 10: The Contextual Transportability of Preferences within Different Models: Marginal Effects

Panel A: Exponential Discounting Model

Cash	Grant	Disc. Fac.(δ)		Risk Aversion (θ)	Income
\$25	\$1000	0.334**		-0.106	-0.135
		(0.074)		(0.113)	(0.086)
\$100	\$1000	0.287**		-0.089	0.043
		(0.065)		(0.083)	(0.076)
\$300	\$1000	0.130**		0.022	0.052
		(0.034)		(0.039)	(0.066)
\$700	\$1000	0.192**		-0.021	0.064
		(0.053)		(0.066)	(0.075)
\$300	\$500	0.377**		-0.162	0.038
		(0.075)		(0.103)	(0.084)
\$300	\$2000	0.175**		0.020	-0.009
		(0.032)		(0.058)	(0.047)
\$300	\$2000	0.474**		-0.020	0.010
		(0.065)		(0.104)	(0.076)

Panel B: Quasi-hyperbolic discount Model

Cash	Grant	Disc. Fac.(δ)	Present bias (β)	Risk Aversion (θ)	Income
\$25	\$1000	0.325**	-0.023	-0.083	-0.135
		(0.074)	(0.043)	(0.114)	(0.087)
\$100	\$1000	0.249**	-0.026	-0.102	0.043
		(0.063)	(0.035)	(0.085)	(0.077)
\$300	\$1000	0.125**	-0.012	-0.002	0.053
		(0.034)	(0.020)	(0.038)	(0.064)
\$700	\$1000	0.163**	-0.006	-0.002	0.064
		(0.051)	(0.029)	(0.068)	(0.077)
\$300	\$500	0.348**	-0.034	-0.144	0.038
		(0.073)	(0.040)	(0.104)	(0.084)
\$300	\$2000	0.173**	0.038	0.033	-0.009
		(0.032)	(0.028)	(0.058)	(0.047)
\$300	\$2000	0.457**	0.035	0.034	0.010
		(0.065)	(0.039)	(0.104)	(0.076)

Panel C: Subjective Risk Model

Cash	Grant	Disc. Fac.(δ)	Payment Prob (p)	Risk Aversion (θ)	Income
\$25	\$1000	-0.033	1.828**	-0.066	-0.108
		(0.054)	(0.574)	(0.082)	(0.087)
\$100	\$1000	0.071	1.240**	-0.087	0.060
		(0.042)	(0.432)	(0.061)	(0.077)
\$300	\$1000	0.037	0.519**	0.012	0.072
		(0.022)	(0.200)	(0.032)	(0.064)
\$700	\$1000	00.028	0.815**	-0.038	0.077
		(0.033)	(0.338)	(0.049)	(0.077)
\$300	\$500	0.048	1.840**	-0.131	0.065
		(0.050)	(0.533)	(0.074)	(0.084)
\$300	\$2000	0.071**	1.060**	-0.011	0.001
		(0.027)	(0.289)	(0.043)	(0.045)
\$300	\$2000	0.139**	3.260**	-0.018	0.035
		(0.051)	(0.581)	(0.075)	(0.073)

Note: (i) Estimates with a "*" are significant at the 1% level. (ii) The marginal effects of income measure the differences in take-up rates between those coming from families earning between \$80,000 and \$100,000 and the reference group (those from families earning \$20,000 or less) ternal validity). While we find evidence of the existence of time-inconsistent preferences, we also stress the empirical relevance of re-interpreting present-bias behavior as the conjunction of individual-specific risk aversion with subjective future payment failure. Finally, among our model representations, long run discount factors, and subjective payment failure are, by far, those with the highest level of transportability (or external validity).

References

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2013). Discounting Behaviour and the Magnitude Effect: Evidence from a Field Experiment in Denmark. *Economica*, 80(320):670–697.

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

Andersen, S., Harrison, G. W., Lau, M. I., and Rutstrom, E. E. (2014). Discounting behavior: A reconsideration. *European Economic Review*, 71:15 – 33.

Andreoni, J. and Sprenger, C. (2012). Estimating Time Preferences from Convex Budgets. *American Economic Review*, 102(7):3333–56.

Bareinboim, E., Lee, S., Honavar, V., and Pearl, J. (2012). Transportability from multiple environments with limited experiments.

Belzil, C., Maurel, A., and Sidibé, M. (2016). Estimating the value of higher education financial aid: Evidence from a field experiment.

Benhabib, J., Bisin, A., and Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, 69(2):205–223.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861–898.

Charness, G., Gneezy, U., and Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, 87(C):43–51.

Cohen, J. D., Ericson, K. M., Laibson, D., and White, J. M. (2016). Measuring time preferences. Working Paper 22455, National Bureau of Economic Research.

Coller, M. and Williams, M. (1999). Eliciting Individual Discount Rates. *Experimental Economics*, 2(2):107–127.

Dekel, E., Lipman, B. L., and Rustichini, A. (2009). Temptation-driven preferences. *The Review of Economic Studies*, 76(3):937–971.

DellaVigna, S. and Malmendier, U. (2006). Paying Not to Go to the Gym. *American Economic Review*, 96(3):694–719.

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

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011). Individual Risk Attitudes: Measurement, Determinants, And Behavioral Consequences. *Journal of the European Economic Association*, 9(3):522–550.

Eckel, C. C. and 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.

Fang, H. and Silverman, D. (2009). Time-Inconsistency And Welfare Program Participation: Evidence From The Nlsy. *International Economic Review*, 50(4):1043–1077.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2):351–401.

Georganas, S., Healy, P. J., and Weber, R. A. (2015). On the persistence of strategic sophistication. *Journal of Economic Theory*, 159(PA):369–400.

Halevy, Y. (2008). Strotz Meets Allais: Diminishing Impatience and the Certainty Effect. *American Economic Review*, 98(3):1145–62.

Holt, C. A. and Laury, S. K. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92(5):1644–1655.

Johnson, C. and Montmarquette, C. (2015). The lack of loan aversion among Canadian high school students. *Canadian Journal of Economics*, 48(2):585–611.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Loewenstein, G. and Prelec, D. (1992). Anomalies in Intertemporal Choice: Evidence and an Interpretation. *The Quarterly Journal of Economics*, 107(2):573–97.

Magnac, T. and Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, 70(2):801–816.

Rosen, H. S. (1985). Housing subsidies: Effects on housing decisions, efficiency, and equity. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of Public Economics*, volume 1 of *Handbook of Public Economics*, chapter 7, pages 375–420. Elsevier.

Wolpin, K. I. and Todd, P. E. (2006). Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility. *American Economic Review*, 96(5):1384–1417.