Estimating the Value of Higher Education Financial Aid: Evidence from a Field Experiment*

Christian Belzil† Arnaud Maurel ‡ Modibo Sidibé §

March 13, 2020

Abstract

Using data from a Canadian field experiment on the financial barriers to higher education, we estimate the distribution of the value of financial aid for prospective students. Our results point out that a considerable share of prospective students are affected by credit constraints. We find that most of the individuals are willing to pay a sizable interest premium above the prevailing market rate for the option to take up a loan, with a median interest rate wedge equal to 6.8 percentage points for a $1,000 loan. The willingness-to-pay for financial aid is highly heterogeneous across students, with preferences and in particular discount factors, playing a key role in accounting for this variation.

Keywords: Higher Education Financing, Time and Risk Preferences, Field experiment.
JEL codes: I22, I23, J24

*First version: May 2016. We would like to thank the Social Research and Demonstration Corporation (SRDC) for allowing us to use the data, which has been developed within the "Willingness to Borrow" project, and may be obtained directly from SRDC. Christian Belzil acknowledges financial support from the ANR "Investissement d'Avenir Labex Ecodec". We thank the Editor, two anonymous referees, Jérôme Adda, Peter Arcidiacono, Jeff Denning, Larry Katz, Claude Montmarquette, Luigi Pistaferri, Basit Zafar, and seminar participants at Cornell, CREST, Paris School of Economics, Toulouse School of Economics, as well as participants at the 2015 World Congress of the Econometric Society (Montreal), the 2017 North American Winter Meetings of the Econometric Society (Chicago), the CEPR Workshop on Financing Human Capital (Sciences Po, Paris, June 2018), and the 2019 SOLE Meetings (Arlington) for useful comments. Brian Clark, Melina Hillion and Jintao Sun provided excellent research assistance.

†CREST, CNRS, Ecole Polytechnique, CIRANO and IZA; christian.belzil@polytechnique.edu
‡Duke University, NBER and IZA; arnaud.maurel@duke.edu
§Duke University and CREST; modibo.sidibe@duke.edu

1 Introduction

Despite inconclusive evidence, many education policies are based on the premise that borrowing constraints preclude students from a modest background to access and persist in higher education. Quantifying the importance of credit constraints in this context is a particularly challenging task, primarily because in most datasets one cannot directly identify the set of students who are constrained. Using data from a unique incentivized artefactual field experiment on the financial barriers to higher education, we propose in this paper a novel empirical strategy based on the estimation of the willingness-to-pay for higher education financial aid.

The field experiment was conducted in Canadian high schools, where students had to make a sequence of choices between a cash payment and various higher education financial aid options: single loan, single grant, and hybrid package composed of a loan and a grant. Loan conditions were similar to the Federal Canadian Student Loan Program. Since the experiment provides education financing ahead of high-school graduation, prospective students facing liquidity constraints are likely to attach a significant value to the opportunity of receiving a loan at the market rate, while those who do not perceive financial constraints should regard those opportunities as redundant.[1] We use this observation to test for the existence of frictions in the market for college loans.

We formulate the decision to accept a financial aid package as an intertemporal problem. The structure of our model may be described as follows. Young individuals, endowed with Constant Relative Risk Aversion (CRRA) preferences, must weigh the increase in utility generated by the acceptance of a cash payment against the expected future gain generated by a specific financial aid package. Although we do not observe the post high-school graduation decisions of agents, we can use information on their revealed preferences to infer the utility gains of accepting different financial aid packages. We derive the willingness-to-pay for financial aid opportunities, and use those estimates to uncover the distribution of the individual-specific additional rates of interest that each individual would be willing to pay to secure financial aid.

To the extent that education financing decisions are made prior to actual college enrollment, the decision to accept a financial aid package depends not only on the perceived magnitude of borrowing constraints, but also on the subjective probability of enrolling in higher education, as well as beliefs over other future outcomes such as probability of dropping-out from college. Understanding financial aid acceptance therefore requires taking into account both heterogeneity in risk aversion and discount factors, as well as financial resources provided by the family. Since the experiment was also designed

[1] The experiment was conducted in several Canadian high schools between 2008 and 2009. Financial aid packages varied from $500 - $4,000 and represented a high fraction of yearly tuition at any of the Canadian. The average tuition was equal to $2,180 for Quebec, $5,667 for Ontario, $3,228 for Saskatchewan, and $5,064 for Manitoba, over the period covered by the experiment. (All amounts are in Canadian dollars, unless noted otherwise.) Similarly, the average US in-state tuition fees charged by public 4-year institutions for 2007-2008 were US $6,200 according to the Bureau of Labor Statistics (Spotlight on Statistics: Back to College, 2010).

to infer fundamental preference parameters for risk and time, it provides a unique opportunity to address this important question and allows us to separate the effects of preferences from other components such as skills, and parental background on the probability to accept financial aid. In the paper, we account for the importance of risk aversion and discount factors, which - thanks to the availability of multiple choices per student - are allowed to be individual-specific, on higher education financing decisions.

The majority of our sample attaches a significant value to the option of accessing higher education loans, and would be willing to trade sizable increases in current consumption in return for the option to take up a college loan at the market interest rate in the near future. Notably, our estimates indicate that the median high school student would be willing to pay an interest rate premium of 6.8 (4) percentage points on top of the prevailing market rate to secure a $1,000 ($ 2,000) loan, consistent with the existence of substantial frictions in the market for college loans. Exploiting the fact that students participated to the experiment at different points in time, we find that, consistent with learning about financial aid opportunities over time, students who participate earlier in their senior year of high school tend to attach larger values to loans. Importantly though, willingness-to-pay for loans remains sizable among those students who are interviewed closer to graduation. The willingness-to-pay for financial aid is also found to be highly heterogeneous across students, while remaining non-negligible for the vast majority of the population. Overall, these results point to the existence of credit constraints which affect a large share of high school students in Canada.

In terms of public policy, our estimates can be more generally interpreted as uncovering the willingness-to-pay for a counterfactual expansion in higher education financial aid offers. Since the set of financial aid offers include grants, our estimates also speak to the value of a significant reduction in university tuition fees. Students attach high values to grants. For instance, our estimation results imply that the median agent would be willing to trade a $450 increase in current consumption against a $1,000 reduction in tuition in the future.

The rest of the paper is organized as follows. In Section 2, we provide an overview of the related literature and detail our contributions. The design of the field experiment and a summary of the data are discussed in Section 3. Sections 4 and 5 describe the model and the econometric specification. Section 6 discusses the identification and the estimation procedure. The estimation results are presented in Section 7. Finally, Section 8 concludes.

2 Background literature and contributions

The existence and the intensity of credit constraints are among the most important issues guiding public policies aimed at stimulating human capital formation such as loans, grants and work study programs (see Lochner and Monge-Naranjo, 2012, 2016, and Heckman and Mosso, 2014 for recent surveys). In the education context, credit constraints denominate any barrier hindering potential

students to finance tuition fees or consumption. At the empirical level, testing for the presence of credit constraints is challenging since borrowing restrictions are not observed in standard data sets. As a consequence, most of the papers have used indirect approaches.

One can distinguish four main strands in this literature. A first set of papers argue that one reason why the estimated returns to schooling using standard instrumental variable techniques may be larger than the OLS estimates is that the subpopulation of compliers tend to be more credit constrained, and thus face larger returns to schooling at the equilibrium (see, e.g., Lang, 1993, and Card, 1995). A second strand of papers test for the importance of credit constraints by estimating short-term effects of parental income on the probability of enrolling in higher education, controlling for long-run factors such as ability (see, e.g., Cameron and Heckman, 1998, Carneiro and Heckman, 2002, and Belley and Lochner, 2007). The third one estimates or calibrates structural models in which credit constraints are explicitly represented, as in Keane and Wolpin (2001), Cameron and Taber (2004), Lochner and Monge-Naranjo (2011), Johnson (2013) and Hai and Heckman (2017). Finally, Cameron and Taber (2004) and Brown, Scholz, and Seshadri (2012) analyze various other testable implications of the existence of credit constraints. With the notable exceptions of Belley and Lochner (2007), Lochner and Monge-Naranjo (2011), Brown et al. (2012) and Hai and Heckman (2017), most of these studies conclude against the existence of a significant role played by credit constraints.

Recently, alternative approaches based on direct measures of credit constraints have been proposed. Although access to these measures provides a clear advantage to researchers, they are typically obtained at the expense of external validity. In addition, while quantifying the overall importance of education financing barriers in the economy requires evaluating their impact *prior* to actual college enrollments, direct measures are generally obtained from a sub-population of individuals who have already enrolled in higher education. A key reference in this literature is Stinebrickner and Stinebrickner (2008), who designed a survey of college students enrolled at Berea College (Kentucky) in order to identify those who are credit constrained and to analyze differences in college drop-out decisions. While a non-trivial fraction of the students declare that they would like to borrow money at a fair interest rate to increase their consumption while in school, the authors conclude that the majority of college attrition is explained by factors other than access to credit.[2]

In this paper, we follow another route and use rich data from a unique college education financing field experiment conducted among Canadian high school students to address this question. By introducing variation in the opportunity costs of accepting various types of higher education financial aid packages, this experiment allows us to estimate the distribution of the monetary values associated with the option to take up college loans at the prevailing market interest rate, and then investigate

[2]Recent work by Delavande and Zafar (2019) investigates the role played by credit constraints in the context of university choice in Pakistan. The authors address the identification issue by directly eliciting from the individuals the university they would have chosen in the (counterfactual) absence of financial constraints, and find that those constraints play an important role.

the existence and intensity of credit market imperfections in this context. Indeed, if agents do not face or anticipate credit constraints, one would expect these values to be small or negligible. On the other hand, large values associated with the option to take up college loans are indicative of credit constraints that are anticipated by the students at the end of high school. That prospective students attach significant values to the option to take up college loans may reflect an adverse effect of credit constraints on future outcomes such as college attendance and graduation, which has often been the focus of the literature. However, students who would be able to enroll in (and graduate from) college in the status quo may also attach significant values to the option to take up loans, reflecting the distortion to consumption profiles induced by credit constraints.[3] Our measure of credit constraints captures these different channels.

Our study is also related to the experimental literature dealing with the estimation of risk and time preferences, including Coller and Williams (1999), Holt and Laury (2002), Andersen, Harrison, Lau, and Rutstrom (2008) and Andreoni and Sprenger (2012) (see Frederick, Loewenstein, and O'Donoghue, 2002, for a survey). In our paper, we estimate non-parametrically the joint distribution of risk aversion parameters and discount factors. This is possible thanks to the large number of choices that are recorded for each individual, which allows us to treat both risk aversion parameters and discount factors as individual fixed-effects. It is worth noting that, in this respect, our approach stands in sharp contrast to most of the dynamic discrete choice literature, where one generally imposes the restriction that all agents share the same discount factor and the same degree of risk aversion.[4] The results obtained in this paper provide clear evidence that discount factors and risk aversion parameters are highly heterogeneous across individuals, and, most importantly, that heterogeneity in preferences plays a central role in explaining the willingness-to-pay for higher education financial aid.

3 Data

The data used in the paper comes from "The Millenium Foundation Field Experiment on Education Financing", which was conducted from October 2008 to March 2009. The initial goal of this experiment was to shed light on the determinants of demand for higher education financial aid separately for different types of financial packages, with a particular focus on the role played by loan aversion in the observed tendency of students to prefer grants over loans.[5] The sample

[3] Related to this point, see Stinebrickner and Stinebrickner (2008), who find that a substantial share of the students enrolled at Berea College are credit constrained, in the sense that they would like to borrow money to increase their consumption in school.

[4] Notable exceptions include Arcidiacono, Sieg, and Sloan (2007) and Brodaty, Gary-Bobo, and Prieto (2014), who allow for heterogeneous discount factors and relative risk aversion, respectively.

[5] See Johnson and Montmarquette (2015), who designed the experiment used in the present paper, for an extensive description of the Millenium Field Experiment. In this paper the authors estimate a reduced-form model of the

consists of 1,248 Canadian full time senior high school students (or students enrolled in CEGEP, the equivalent of senior year of high school in Quebec), aged between 16 and 18 years old from both urban and rural sites across Canada.[6] The experiment was funded by the Canada Millennium Foundation, a public enterprise created by the Canadian federal government, and was carried jointly by The Social Research and Demonstration Corporation (SRDC, Ottawa, Canada) and the Centre Interuniversitaire de Recherche en Analyse des Organisations (CIRANO, Montreal, Canada).

3.1 The experiment

The experiment was conducted using pen and paper choice booklets as well as simple random sampling devices like bingo balls and dice. Project cost considerations suggested that participants be drawn from locations with convenient travel connections from the SRDC Ottawa and CIRANO Montreal offices. Manitoba, Saskatchewan, Ontario and Quebec were the selected provinces. The implementation team was able to carry out work in urban and rural schools in each of the four provinces.[7] Table A3 in Appendix B provides a descriptive overview of our sample.[8]

The experiment consists of three parts. First, students must answer a set of questions aimed at measuring their rate of time preference. In those questions, individuals are offered a choice between two payments of different values to be made at different points in time. Our analysis uses the 24 questions that provide a cash payment within a day or a week (see Table A1 in Appendix B). A second set of questions relate to the measurement of risk attitudes. Specifically, students are presented with a sequence of 55 binary choices between two lotteries in which risk is objectively stated (see Table A2). Finally, the third group of questions is a sequence of choices between a cash payment to be paid within one week from the day the experiment was carried, and the option to take up a specific financial aid package covering future educational expenses. These questions play a key role in our empirical strategy. Should the student decide to exercise this option, the financial

decision to take up loans and conclude against the existence of sizable loan aversion.

[6] In practice, 53% (47%) of the sample was collected in the Fall (Spring) semester.

[7] Experimental staff was granted access to the high schools and cooperated with student services staff to recruit and schedule senior students. Sessions took place during school hours (over two days). Participation to a Web survey and parental consent were required to participate in the experimental session. The experimental sessions were held in controlled environments including classrooms, libraries, career counseling rooms, activity rooms and auditoriums. All rooms were held on the campus where the student attended classes. The planned optimal number of participants per session was between 20 and 25 allowing the entire urban subject pool to be contained in 50 sessions. A total of 75 sessions were conducted with 50 as the maximum number of participants in any session. For showing up on time, each participant received a $20 fee. Supplementary information regarding aspects of the experiment that we do not model may be found in Johnson and Montmarquette (2015).

[8] Although the first objective of the designers of the experiment was not to achieve national representativity, comparison with nationally representative data from the Youth In Transition Survey (YITS) suggests that the resulting sample is largely representative of the subpopulation of Canadian youths who have obtained at least a high school degree. See Belley, Frenette, and Lochner (2014) who provide a detailed discussion of the YITS data in the context of post-secondary attendance decisions in Canada, and Hansen and Liu (2013) who report the characteristics of individuals from the YITS sample separately by schooling level.

aid package is to be paid conditional on enrolling in a full-time program at any higher education institution in the country (within two years).[9] Importantly, choices were incentivized as students were paid for one randomly drawn decision at the end of the session.

Overall, three types of financial aid packages were offered, namely grants, loans and hybrid loans which combine a loan and a grant. We use a total of 17 financial aid decisions for each respondent, with 5 choices with a single loan offer, 7 choices with a single grant offer, and 5 hybrid loan offers. These decisions are summarized in Table 1 below. Participants were told that grants and loans would be disbursed upon enrollment in a higher education institution for learning or training full time within two years from the date of experiment participation.[10] Grants and loans were defined as follows:

- Grants: cover direct and indirect costs related to the learning activity. For tuition fees, payments will be made directly to the education institution. Receipts will be required for the reimbursement of other costs.

- Loans: are available up to two years from the date of the experiment. The loans are repayable upon completion of, or dropping out from the program of study. The interest rate, which is the same as the one offered by the Canadian Federal Student Assistance program, is floating and is set at the prime rate (3.2% on average over the period of interest) plus 2.5%.

Cash alternatives varied from $25 to $700, while grants and loans varied from $400 to $4,000. The variations in cash amounts and in the nature and the size of financial aid packages play a crucial role in our analysis.

[9] Over this period in Canada, the transition rate from high school to higher education was around 85% (see Belley et al., 2014).

[10] Specifically, within two years of the experiment, participants had to contact the project manager (SRDC in Ottawa) and show an official proof of enrollment from the Registrar of the university or college. This implied that to get financial aid, the students needed to have paid the minimum tuition (usually one semester) required by the institution to be granted a "registered" status.

Table 1: Financial aid questions

Type of package	Choices	Cash	Grant	Loan	Aid Take-up
	1	25	0	2,000	0.458
	2	300	0	2,000	0.171
Single Loans	3	700	0	2,000	0.051
	4	300	0	1,000	0.109
	5	300	0	4,000	0.283
	6	25	1,000	1,000	0.834
	7	300	1,000	1,000	0.637
Standard Hybrid	8	700	1,000	1,000	0.389
	9	300	400	400	0.287
	10	300	2,000	2,000	0.727
	11	25	1,000	0	0.886
	12	100	1,000	0	0.826
	13	300	1,000	0	0.686
Single Grants	14	700	1,000	0	0.412
	15	300	500	0	0.384
	16	300	2,000	0	0.764
	17	300	4,000	0	0.835

Notes: i) Amounts are in Canadian dollars, ii) Source: SRDC-CIRANO Field Experiment on Education Financing.

At the outset, it should be clear that these amounts are quite significant in the Canadian context. For instance, over the period considered, a grant of $2,000 in 2008 would have covered 65% of yearly fee at University of Western Ontario and Queen's University, and almost 100% at McGill University and Université de Montréal. Importantly for our empirical strategy, loans in the experiment were offered at the same interest rate as in the Federal Canadian Student Loan Program. At the time of the experiment, this rate, which we sometimes simply refer to in the paper as the "market rate", was equal to 5.7%.[11] In Appendix B we provide additional details about how the grants and loans were administered as part of the experiment. Overall, this experiment was a major undertaking. Put in 2015 Canadian dollars, the total cost of the experiment was around $1,000,000.

Finally, over the day of the experiment, a numeracy test provided by the Center for Education Statistics was administered to all students. The test was based on the numerical component of the

[11] The interest rates of the loans offered through the experiment were also similar to the rates offered by Canadian Charter banks such as the Royal Bank, the Toronto Dominium Bank and the Bank of Montreal, within their private education loan programs. For students enrolled in specific fields such as Medicine or Engineering, private education loans were sometimes offered at lower rates than federal loans (Annual Report about the Canadian Student Loans Program, Human Resources and Skills Development Canada, 2010-2011). See Appendix A for an overview of higher education financial aid in Canada.

International Adult Literacy and Skills Survey project undertaken by numerous OECD countries between 1995 and 2005. The questions are meant to capture the capacity to perform numerical calculations (Murray, Clermont, and Binkley, 2005). Students received a score between 0 and 500, which we use in our analysis as a measure of cognitive ability.[12]

3.2 Describing the take-up rates

To describe the sensitivity of individual take-up to financial aid structure and cash payments, we plot below the take-up rates associated with various combinations of grants and loans, against specific cash payments. In Figure 1, the take-up rates are those obtained when the alternative cash payment was \$25 and \$700, respectively.[13] As expected, the take-up rates are inversely related to the amount of cash payment. For instance, 90% of individuals opted for a \$1,000 grant when offered a \$25 cash payment, while only 40% opted for a \$1,000 grant when offered a cash payment of \$700. Those differences in take-up rates are very similar for the three financial aid packages considered in this figure (\$1,000 grant, hybrid \$1,000 loan combined with \$1,000 grant, and \$1,000 loan). Figure 2 reports the take-up rates for various sizes and types of financial aid packages and for a \$300 cash payment. This figure shows that, for all three types of packages, take-up rates increase with the size of the financial aid. However, the increase in take-up rates is relatively modest after \$1,000. Figure 2 shows that the marginal increase in take-up rates for additional financing is relatively small, with take-up rates respectively equal to 76% for a \$2,000 grant and 84% for a \$4,000 grant, against 68% for a \$1,000 grant and 39% for a \$500 grant. An explanation for those take-up rates being significantly below 1 even for a \$4,000 grant is that not all individuals in our sample intend to enroll in higher education, and as a consequence, some individuals will always value more current consumption over higher education financing opportunities. A similar pattern can be observed for single loan offers, as the take-up rate for \$1,000 loan is 11%, compared to 17% and 24% for \$2,000 and \$4,000 loans, respectively.

Figures 1 and 2 also show that take-up rates for a grant of value x are very similar, and in fact always slightly higher, than for a hybrid financial package offering an additional loan of value x. It is important to note that this pattern does not imply that students do not attach positive values to the option to take up a loan. Instead, this likely reflects the fact that taking up a hybrid package in practice entails taking up both the loan and the grant associated with it, and thus subsequently paying the loan off with interest. Depending on the expected future consumption profile and intensity of credit constraints, individuals may rationally attach a higher value to a grant offer than to a hybrid package adding a loan to the grant. That the take-up rates for single loan offers increase with the loan amounts, while take-up rates for hybrid packages are lower than for

[12] In our empirical application, individual numeracy test score has been rescaled to have variance one, with a mean equal to 5.1.

[13] L stands for loan, G for grant and H for hybrid.

Figure 1: Take-up rates against $25 and $700 cash

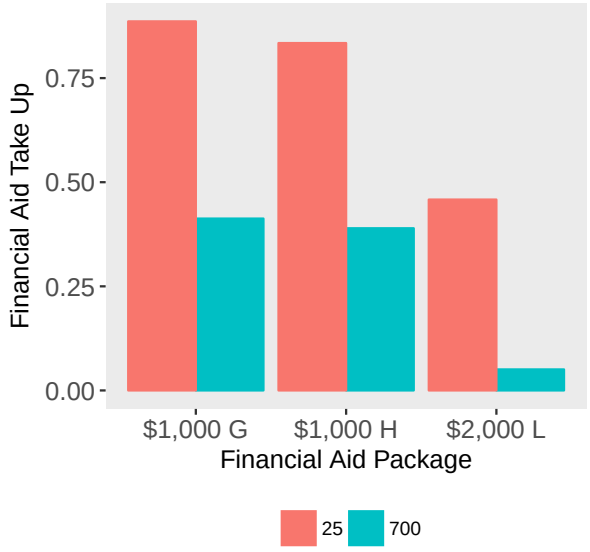

grants, further points to the future value of financial aid being non-linear in post-college repayment amounts. The choice model we consider in the next section allows for such non-linearities.

4 The model

We present a simple choice model which, combined with our experimental data, allow us to quantify the willingness-to-pay for financial aid opportunities. Namely, we assume that preferences over consumption are represented by a CRRA utility function. For each student i, the utility of consumption c is given by:

$$u(c, \theta_i) = \begin{cases} \dfrac{c^{1-\theta_i}}{1-\theta_i} & \text{if } \theta_i \neq 1 \\ \log(c), & \text{if } \theta_i = 1 \end{cases} \tag{1}$$

where θ_i denotes the individual-specific risk aversion parameter.

The large number of questions available from the experiment (96 in total) allows us to treat relative risk aversion rate as well as the annual discount factors as individual-specific parameters. Because individuals differ not only in their preferences, but also in their financial resources, we further allow for individual specific consumption endowments. In the rest of the section we follow the structure of the experiment and present the choice equations separately for each group of questions.

Figure 2: Take-up rates against $300 cash

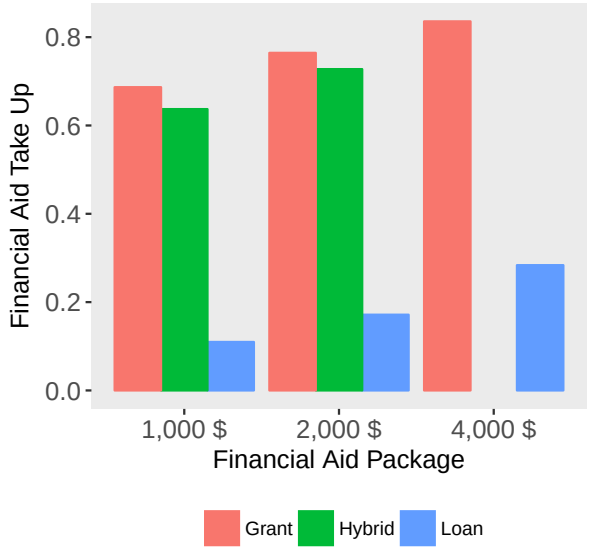

4.1 Time preference

A first set of questions (indexed by $q = 1, \dots 24$) is devoted to measuring the subjective discount rate and consists of a sequence of choices between two alternatives: an early cash payment denoted by a_{0q} to be paid t_{0q} months from now, and another cash payment denoted by a_{1q}, and to be paid t_{1q} months from now. We denote by Y_{iq} a dummy variable which is equal to 1 if individual i chooses the early cash payment. Questions differ with respect to the amount of the cash payment, and the periods at which the earlier or later cash payments would be paid. As is standard in the experimental literature on the estimation of time and risk preferences, we assume in the following that individuals immediately consume cash payments upon reception (see, e.g., Andersen et al., 2008).

We express individual i's probability to choose consumption at period t_{0q} versus period t_{1q} using a simple two-period consumption model:

$$\Pr(Y_{iq} = 1) = \Pr\Big\{ \beta_i(t_{0q})[u(c_i + a_{0q}, \theta_i) - u(c_i, \theta_i)] + \beta_i(t_{1q})[u(c_i, \theta_i) - u(c_i + a_{1q}, \theta_i)] + \varepsilon_{iq} > 0 \Big\} \tag{2}$$

where ε_{iq} is an idiosyncratic preference shock and c_i the individual background consumption. The distributional assumptions and specification of the background consumption are discussed in Section 5. Finally, the individual discount rate applied for a payment to be received t_q months from

now is assumed to be given by:

$$\beta_i(t_q) = \frac{1}{1 + \frac{t_q}{12} \cdot r_i}$$

where r_i denotes the individual annual subjective interest rate.

4.2 Risk aversion

A second set of questions relate to the measurement of the degree of relative risk aversion. In each of these questions, individuals are offered to choose between two alternatives, namely a lottery offering a^1_{0q} with probability p_q and a^2_{0q} with probability $1-p_q$, and another lottery offering a^1_{1q} with probability p_q and a^2_{1q} with probability $1-p_q$. Questions differ according to the state contingent cash payments $(a^1_{0q}, a^2_{0q}, a^1_{1q}, a^2_{1q})$ and probabilities (p_q, $1-p_q$). The generic choice probability for the first alternative is given by:

$$\Pr(Y_{iq}=1) = \Pr\Bigg\{ p_q \cdot \Big(u(c_i + a^1_{0q}, \theta_i) - u(c_i + a^1_{1q}, \theta_i)\Big) \\ +(1-p_q) \cdot \Big(u(c_i + a^2_{0q}, \theta_i) - u(c_i + a^2_{1q}, \theta_i)\Big) + \varepsilon_{iq} > 0 \Bigg\} \tag{3}$$

where ε_{iq} is an idiosyncratic preference shock, and $q = 25, \ldots 79$.

4.3 Choices between consumption and education financing

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

In order to interpret individual choices between cash payments and educational financial aid, we specify a stylized two-period model. Period 0 refers to the time when high school students are asked to choose between consumption and an education financing package, while period 1 refers to the residual life-cycle starting from high school graduation. For each choice, individuals must decide between a cash payment and a financial aid package which is to be paid conditional on higher education enrollment in period 1.

This structure of the model is as follows. Let a_{0q} and a_{1q} be the cash payment and financial aid transfer in question q. The potential financial aid offer (a_{1q}) offered in period 0 and to be paid at the beginning of period 1 is characterized by a two-element vector, denoted by $(g_q, \ell_q)'$ where g_q is the amount of the grant and ℓ_q the amount of the loan. The choice variable Y_{iq} is equal to 1 when an individual chooses financial aid a_{1q}, and 0 if she accepts consumption a_{0q}. As a consequence, the initial period consumption c_{0iq} is given by

$$c_{0iq} = c_i + a_{0q} \cdot (1 - Y_{iq})$$

The period 0 utility, denoted by $W(.)$, of accepting or rejecting the financial aid is then given by:

$$W(Y_{iq} = 1) = u(c_i, \theta_i) + \varepsilon_{1iq} \quad (4)$$
$$W(Y_{iq} = 0) = u(c_i + a_{0q}, \theta_i) + \varepsilon_{0iq} \quad (5)$$

where ε_{0iq} and ε_{1iq} represent choice specific preference shocks. Implicit here is the assumption that individuals immediately consume the cash payment, thus ruling out situations where high school students would instead save some of the cash transfers for college. In other words, we assume that, at the time of the experiment (end of high school), agents are borrowing constrained and, as a consequence live hand-to-mouth.[14] In Section 7.3.2 we show that our conclusions remain unchanged after excluding the questions with the highest cash payment ($700) for which this assumption may seem more restrictive.

Because we do not have data on actual choices exercised in subsequent periods, we focus on the estimation of the difference between future components of the utilities of accepting and rejecting financial aid. The intertemporal utilities of accepting and rejecting financial aid are given by $u(c_i, \theta_i) + \varepsilon_{1iq} + \beta_i EV_{iq}(Y_{iq} = 1)$ and $u(c_i + a_{0q}, \theta_i) + \varepsilon_{0iq} + \beta_i EV_{iq}(Y_{iq} = 0)$, respectively, where the future components $EV_{iq}(Y_{iq} = 1)$ and $EV_{iq}(Y_{iq} = 0)$ are the expected lifetime utilities associated with each choice.

In practice, expected future value terms associated with each alternative $Y_{iq} \in \{0, 1\}$ depend on individual beliefs about a range of future outcomes, some of them being alternative specific. Notably, those beliefs include the subjective probabilities of enrolling in higher education conditional on receiving the financial aid offer a_{1q}, and conditional on not receiving the financial aid offer, for the alternatives $Y_{iq} = 1$ and $Y_{iq} = 0$ respectively. The expected future value terms in principle also depend on the perceived availability of higher education financing opportunities outside of the experiment, as well as possibly on beliefs over other future outcomes such as the probability of dropping-out from college. Any crowding-out of parental transfers by loans or grants (Lochner, Stinebrickner, and Suleymanoglu, 2018; Hotz, Wiemers, Rasmussen, and Maxwell Koegel, 2018), in as much as it is anticipated by the students at the time of the experiment, will also be reflected in the future value of financial aid. As those outcomes are not observed in the data, we treat those future components as unknown functions of parental socio-economic background and individual skills, as well as risk and time preferences. A key advantage of this approach relative to a

[14] The assumption that agents live hand-to-mouth is common in dynamic structural models, and has been invoked in a variety of contexts. See, among others, Humphries (2018), Pavan (2011), Carneiro, Hansen, and Heckman (2003), Rust and Phelan (1997), for models of self-employment, job search and career choice, college enrollment, and joint labor supply and Social Security acceptance, respectively. Relaxing this assumption would require introducing an asset market, which would complicate the model (and estimation) substantially.

more standard dynamic discrete choice model is that it avoids strong and untestable restrictions regarding the beliefs of the individuals and their evolution over the life cycle.[15]

Under these assumptions, agent i's probability to accept the financial aid package proposed at question q is given by

$$\Pr(Y_{iq} = 1) = \Pr\left\{u(c_i, \theta_i) - u(c_i + a_{0q}, \theta_i) + \beta_i \cdot \psi_{iq} + \varepsilon_{iq} > 0\right\} \tag{6}$$

where $\psi_{iq} = EV_{iq}(Y_{iq} = 1) - EV_{iq}(Y_{iq} = 0)$ denotes the expected future utility gain associated with accepting financial aid, and $\varepsilon_{iq} = \varepsilon_{1iq} - \varepsilon_{0iq}$.

5 Specification

In the following, we assume that the idiosyncratic shocks are independent across individuals and questions, and identically distributed within each of the three groups of questions following a normal distribution with mean zero.[16]

5.1 Background consumption

Until now, we have simplified the exposition by considering for each individual a unique background consumption variable c_i. However, in order to capture the differences in stakes between the discount rate and risk aversion questions, on the one hand, and the questions about financial aid, on the other hand, we allow the background consumption levels, and therefore also the marginal utility of consumption to vary across groups of questions.[17] Specifically, for each group of questions k (where $k \in \{1, 2, 3\}$ indexes the questions related to time preference, risk aversion and education financing, respectively), we denote the individual background consumption by c_i^k, which is assumed to depend on a vector of individual and family background characteristics denoted by Z_i:

$$c_i^k = \exp(\gamma_k Z_i) \tag{7}$$

[15] An alternative approach would be to elicit beliefs about counterfactual future schooling choices and consumption paths. In the higher education context, notable references include Arcidiacono, Hotz, and Kang (2012), Stinebrickner and Stinebrickner (2014), Wiswall and Zafar (2015a, 2015b) and Zafar (2013) who use subjective expectations data to examine the determinants of college major choice.

[16] Although the error terms are assumed to be independent across questions, individual choices are correlated across questions through the individual-specific risk and time preference parameters, and of course via their observed background characteristics too. The assumption that the idiosyncratic shocks are i.i.d. across questions is standard in the experimental literature (Andersen et al., 2008). Given the large number of questions, relaxing this assumption would very significantly add to the computational burden.

[17] As such, background consumption is best interpreted as an individual-specific parameter that is picking up the monetary stakes associated with the different groups of questions. In practice, allowing c_i to be specific to each of the three groups of questions proved important in fitting the financial aid choice data.

The vector Z_i includes an intercept and a set of individual characteristics, namely gender, parental income (20-40K, 40-60K, 60-80K, 80-100K and more than 100K Canadian dollars per year), parental education (high school, vocational college and college) of the respondent of the parental survey, Canadian citizenship, place of residence (Quebec, Ontario, Saskatchewan, and a dummy for rural location) and family composition (existence of siblings younger or older than 18).[18] This flexible specification allows each of these characteristics to have different weights within each group of questions.

5.2 Approximation of the future component

The expected future utility gain of accepting financial aid (ψ_{iq}) is a key component of the model. Standard structural dynamic estimation would require to posit a parametric form of the utility of potential future schooling choices as well as all the probability distributions that characterize the subjective beliefs of agents about future returns to education and experience. In this paper, we take another route and assume instead that the future component may be approximated by a parametric function that depends on some observable characteristics, namely loan and grant amounts, skills (as measured by numeracy test scores), geographical location, parental income and education. Importantly, the future component also depends on two unobserved individual characteristics, namely the discount factor and risk aversion parameter, which are treated as individual fixed effects in our analysis.[19] At a high level, this is analogous to a factor model where we have access to noisy measurements of two correlated factors, namely time and risk preferences. We estimate this function using a flexible polynomial incorporating those variables in level and square as well as various interactions.[20]

[18] There are 269 individuals with missing income data, and 146 individuals with missing education data. We use a Gibbs sampling algorithm to impute those missing information using age, gender, location, income, education, citizenship variables in the full sample.

[19] One potential implication of not allowing for additional sources of unobserved heterogeneity in the future component is the overestimation of the variances of the idiosyncratic shocks. In our baseline specification, these shocks are excluded from ψ_{iq}. The standard deviation of the preference shock for financial aid questions is 0.131 to be compared to 1.99 for discounting questions and 1.74 for risk aversion decisions. Our conclusions remain unchanged when we include instead the idiosyncratic shocks in the future value component (see Section H.1).

[20] The exact specification is given in Appendix D. Interaction between risk aversion and discount factor, squared terms in those preference parameters, and dummies for the existence of siblings younger or older than 18 are excluded from this specification as they were not statistically significant when we allowed them to enter the future component. Our approach is in line with Geweke and Keane (2000), who advocate the estimation of the future component of the value function using a polynomial which records the movement of state variables implied by the law of motion. However, in contrast with Geweke and Keane (2000), we account for heterogeneity in individual preferences and incorporate time and risk preference parameters in the polynomial.

6 Identification and estimation

In this section we first discuss the sources of identification of the key parameters of the model, before turning to the estimation procedure.

6.1 Identification

We start by informally discussing the sources of identification of the individual time preference and risk aversion parameters, before turning to the parameters of the expected future utility gain of accepting financial aid. We highlight below the main features of the experiment and of the model that are key to identification. While the experiment does provide us with separate sources of identification for those different sets of parameters, (point-)identification remains parametric in the sense that it relies on the distributional assumptions made on the idiosyncratic shocks.[21]

First, the risk aversion parameters are primarily identified from the sequence of choices described in Section 4.2 between two lotteries that differ in the cash payments and the probabilities associated with each payment. Importantly, variation across individuals in time preferences does not confound the identification of the risk aversion parameters here since all of the state-contingent cash payments are immediate.

The main sources of identification of the individual discount rates are the choices between earlier and later cash payments that are described in Section 4.1. However, given that in our model individuals are endowed with non-linear CRRA preferences over consumption, part of this variation may of course also reflect heterogeneity across individuals in risk aversion. This highlights the importance of having access to choices that only depend on risk aversion, and not on discount factors.

Finally, having identified the time and risk preference parameters from the choices described above, the expected future utility gains of accepting financial aid (versus cash payment) are identified from the choices between immediate cash payment and education financial aid that are modeled in Section 4.3. Specifically, under our parametric assumptions, the expected future utility gains associated with the various types of financial aid are identified from the variation in take-up rates across discount rates. It is interesting to note that, at this stage, individual-specific discount rates – which are identified in the previous step from the choices between earlier and later cash payments – effectively play the role of an exclusion restriction, in that they only affect the choices through the future component of the utility. As such, aside from being of interest in its own right, the distribution of discount rates plays an important role in the identification of the willingness-to-pay for different financial aid packages.

[21] This should not come as a surprise since binary choice models are generally not point-identified without making distributional assumptions on the error term. Note that this is true even if the parameters of interest are not individual-specific (see, e.g., Magnac and Maurin, 2008).

6.2 Estimation

We estimate the model using a stepwise maximum likelihood procedure. We resort to a sequential approach to keep the estimation tractable despite the large dimensionality of the optimization problem.[22] The first step consists of estimating the risk aversion parameters $(\theta_i)_i$ and the background consumption parameters $(c_i^1)_i$ using the subset of questions designed to elicit risk preferences. Holding those parameters fixed, we then estimate in a second step the individual subjective interest rates $(r_i)_i$ and the background consumption parameters $(c_i^2)_i$ using the time preference questions. Finally, we estimate the future component of accepting the different types of financial aid $(\psi_{iq})_{i,q}$ and the background consumption parameters $(c_i^3)_i$ using the school financing decisions.[23] Standard errors are derived using bootstrap with 500 replications.

7 Results

We present the estimation results as follows. The first part is devoted to the distributions of risk aversion and discount factors, and of the background consumption levels. The second part discusses the estimates of the expected future utility gain of accepting financial aid, for various types of financial packages. In the third part, we present the model fit. We introduce and discuss the willingness-to-pay for financial aid in Part 4, and then study the implications of these results in terms of credit constraints. Finally, we investigate the relative importance of family background characteristics, individual skills as well as time and risk preferences as determinants of the willingness-to-pay for financial aid.[24]

7.1 The distributions of risk aversion, discount factors and background consumption

The empirical distributions of the estimated degrees of relative risk aversion, $(\theta_i)_i$, and discount factors, $(\beta_i)_i$, are reported in Figure 3, while Table A9 describes the joint distribution of the estimated relative risk aversion parameters and discount factors. Starting with risk aversion, 97.2% of the estimated risk aversion parameters are significant at 5%, thus leading to the rejection of risk neutrality ($\theta = 0$) for the vast majority of the individuals in the sample. The distribution

[22] Note that the properties of the CRRA function imply that the likelihood function is not continuous. As such, standard quasi-newton optimization techniques can not be used to obtain asymptotically efficient estimates using the sequential estimates as starting values.

[23] Our framework corresponds to a binary choice panel data model with fixed effects, where the longitudinal dimension is given by the various questions q. It is therefore in principle subject to the incidental parameter problem (Neyman and Scott, 1948). However, in practice we expect the incidental parameter bias to be negligibly small here given the large number of questions that are used in the estimation (96 overall for each individual).

[24] Estimation results for the background consumption parameters (γ_3) associated with the financial aid questions are reported in Appendix E. Estimation results for the other background consumption as well as the future component parameters are available from the authors upon request.

of risk aversion is skewed to the right and is characterized by a substantial degree of dispersion, as indicated by the inter-decile range (0.40 for the first decile and 0.85 for the last decile). Risk loving behavior ($\theta < 0$), however, is limited to 16 individuals out of 1,248. Regarding the external validity of these results, it is worth noting that the median of the estimated risk aversion parameters (0.64) fits in the range of the relative risk aversion parameters that have been estimated, using observational data, in the literature on dynamic schooling choices. For instance, Keane and Wolpin (2001) estimate a smaller risk aversion coefficient (0.49), while Belzil and Hansen (2004) and Sauer (2004) find larger degrees of risk aversion (0.93 and 0.77, respectively). Particularly relevant for us is the paper by Brodaty et al. (2014), who, unlike the previous studies, estimate a dynamic model of schooling decisions that allows for heterogeneous degrees of relative risk aversion across individuals. In their paper, the estimated risk aversion coefficients range between 0.6 and 0.9.

Figure 3: Histogram of the risk aversion and discount factors

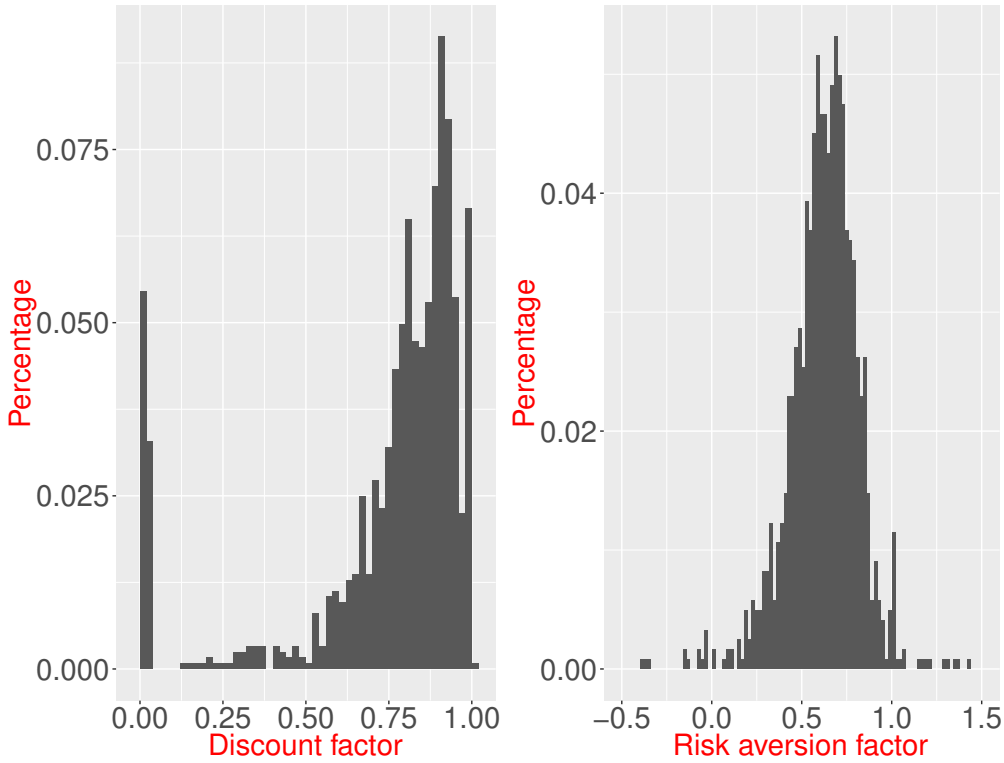

Notes: (i) The range of the risk factor has been restricted from -0.5 to 1.5 to improve display, excluding a total of 24 individuals. (ii) The binwidth has been set to 0.02.

Turning to the time preferences, the discount factors are also found to be quite dispersed. The distribution is skewed to the left with its median (0.83) being larger than its mean (0.75). The empirical distribution of discount factors is essentially bimodal, with around 10% of our sample

Table 2: Background consumption

	Discount Questions	Risk Questions	Financial Questions
Mean	3.21	0.58	237.30
SD	36.37	1.02	274.75

having an estimated discount rate lower than 0.33, while another 40% of the sample has discount rate higher than 0.87. This points to the co-existence of a myopic sub-population and a set of forward-looking individuals endowed with large discount factors of about 0.9. In Appendix H.1, we show, however, that our main results in terms of willingness-to-pay for financial aid are not driven by respondents who have very low discount rates. Comparing our discount factor estimates with those obtained in earlier studies, it is interesting to note that our mean value matches that of Andersen et al. (2008) who report, using data from an artefactual field experiment conducted in Denmark among adults aged 19 to 75, an estimate of 0.75 (under a risk neutrality assumption), while being lower than the value estimated under risk aversion (0.9).[25]

Our data also allows us to examine the joint distribution of risk aversion and discount factors. Table A9 reports a negative correlation (-0.14, significant at 1%) between the discount rate and relative risk aversion. In the experimental economics literature, risk and time preferences are generally not elicited jointly and for this reason there exists only few estimates of the correlation between risk aversion and discount factors. A notable exception is Andersen et al. (2008), who discuss the bias affecting discount factor estimates when individuals are erroneously assumed to be risk neutral, and show that the joint elicitation of risk and time preferences results in lower discount rates estimates. Andersen et al. (2008) consider the joint distribution of risk aversion and discount rates using a parametric model in which both risk aversion and discount rates depend on observed heterogeneity and an orthogonal unobserved heterogeneity term. Consistent with our results, they report a weak positive correlation between risk aversion and impatience.

Finally, in our model, the level of background consumption is allowed to vary by family background, and across groups of questions as well. This specification, which allows us to capture the important heterogeneity in stakes associated with the different types of questions - in particular financial aid questions versus questions pertaining to risk and time preferences - proved important in fitting the distribution of individual choices. The estimates reported in Table A11 indicate that the consumption levels that are used to evaluate cash payments do differ across questions. In particular,

[25]Some of the papers estimating dynamic models of schooling decisions also attempt to estimate the discount factor. Notably, Keane and Wolpin (1997) estimate for their baseline model a discount factor (0.85) that is very close to the median of our estimated distribution (0.83). See Magnac and Thesmar (2002) who provide sufficient conditions under which the discount factor is identified in the context of dynamic discrete choice models.

the background consumption levels for the financial aid questions are, on average, much higher than for the other questions, while being also highly dispersed across individuals.

Having estimated the time and risk preferences for each individual in the sample, it is interesting to examine whether these parameters could be predicted by standard background characteristics. Table A10 in Appendix E reports the results from a linear regression of the estimated individual-specific degrees of relative risk aversion and discount factors on a set of socio-economic background and demographic characteristics. The main takeaway from this table is that those characteristics only account for a very small fraction of the variation in risk aversion parameters and discount factors, the R^2's for both regressions being low (0.02 and 0.09, respectively). Nonetheless, it is worth noting that females are significantly less forward looking and more risk averse than males. Immigrants also tend to be less forward looking, while children of college educated parents have a lower degree of risk aversion, although the coefficient is only significant at the 10% level. Finally, residents of Saskatchewan, which is the poorest region in the sample, have significantly lower discount rates. At any rate, these results show that discount rates and risk aversion parameters are economic primitives which are mostly left unexplained by standard background and demographic characteristics.

7.2 The expected future gain of financial aid

In this subsection, we focus on the relevant decision variables between cash payment and financial aid. Specifically, Table 3 compares the distribution of the expected discounted future utility gain against the period 0 utility of accepting a cash payment, for various types of financial aid packages and cash payments. Period 0 utility gains of accepting a cash payment (denoted by $\Delta(.)$) are evaluated as the difference in the initial period utility of consumption evaluated at the reference consumption level (c_i) plus a cash payment, and the initial period utility evaluated at the reference consumption level.

First, as expected, for a given amount of financial aid, the median value of a grant exceeds the value of a loan. More generally, the distribution of the discounted expected future utility gains for a grant first-order stochastically dominates that of the expected future gains for a loan. Second, the utility gain associated with a grant increases with its size, with the distribution of the utility gains of a \$2,000 grant ($\beta\psi(g_{2000})$) dominating that of a \$1,000 grant ($\beta\psi(g_{1000})$). A similar pattern holds for loans. This result shows that, in most parts of the distribution, the interest costs are not high enough to make a \$2,000 loan less attractive than a \$1,000 loan. Finally, the discounted expected future utility gains of the financial aid packages vary a lot across individuals, with the ratio of interdecile range over the median ranging from 1.1 (for a \$2,000 grant) to as much as 2.9 (for a \$1,000 loan).

We now examine the utility gains of accepting a cash payment. The distributions of the utility

Table 3: Utilities of consumption and expected discounted future gain of financial aid

	$\Delta(c_{100})$	$\Delta(c_{300})$	$\Delta(c_{700})$	$\beta\psi(\ell_{1000})$	$\beta\psi(\ell_{2000})$	$\beta\psi(g_{1000})$	$\beta\psi(g_{2000})$
Min	0	0	0	0	0	0	0
1st Dec.	32	83	157	2.8	2.8	47	96
1st Quart.	48	122	223	15	19	148	237
40th Pct.	58	144	260	29	36	186	282
Median	66	156	280	40	50	206	305
Mean	74	167	303	62	68	210	301
SD	144	152	282	144	142	167	190
60th Pct.	73	173	305	55	64	227	330
3rd Quart.	84	198	345	77	86	263	370
9th Dec.	103	229	407	117	122	321	438
Max	4281	4282	7599	3048	3014	3214	3448

Notes: (i) $\Delta(c_x) = u(c_i + x) - u(c_i)$, (ii) $\beta\psi(\ell_{1000})$ (resp. $\beta\psi(\ell_{2000})$) denotes the expected discounted future utility gain associated with a \$1,000 (resp. \$2,000) loan, (iii) $\beta\psi(g_{1000})$ (resp. $\beta\psi(g_{2000})$) denotes the expected discounted future utility gain associated with a \$1,000 (resp. \$2,000) grant.

gains associated with a \$300 ($\Delta(c_{300})$) and \$700 ($\Delta(c_{700})$) cash payments dominate that of both types of loans. This is also true throughout most of the distribution of the utility gains associated with a smaller \$100 cash payment ($\Delta(c_{100})$). On the other hand, the sign of the value of immediate cash versus grant varies across the distribution, with the median utility gain of accepting a \$700 cash payment ranging between the median gain of a \$1,000 grant and that of a \$2,000 grant.

Finally, recall that the expected future utility gains associated with the various types of financial aid are identified from the variation in take-up rates across discount rates. That the estimated future utility gains are sizable - for instance for the case of a \$2,000 grant, generally larger in magnitude than the utility gains of accepting a \$300 cash payment - is interesting in that respect. Indeed, this points to the external validity of discount rates which are identified primarily from the sequence of lower-stakes decisions involving payments at different points in time, but are nevertheless predictive of financial aid decisions.

Taken together, these results already provide suggestive evidence that the willingness-to-pay for financial aid packages is heterogeneous across high school students. However, our estimates allow us to go beyond the marginal distributions of utility gains associated with cash payments and financial aid, and directly compute individual-specific willingnesses-to-pay for the different types of financial aid packages that are proposed in the experiment. We report and discuss the estimated distributions of the willingness-to-pay for various types of financial aid packages in Section 7.3.

7.3 Borrowing constraints and the value of financial aid

7.3.1 Measurement

In this section, we combine the observed choices between cash transfers and financial aid packages with our model to evaluate the monetary values associated with the option to take up the various types of financial aid. Before doing so, it is important to examine the ability of our model to fit the financial aid decisions from the experiment. Table A12 in Appendix F reports the empirical frequencies of financial aid acceptance and the predicted probabilities, for each of the 17 financial aid choices to exercise. Our model generally fits the data reasonably well, with a couple of exceptions. Most notably, we underestimate the take-up rates of a \$1,000 grant and of the hybrid package (\$1,000 loan and \$1,000 grant) against a \$700 cash payment, while we overestimate the take-up rates for a \$2,000 loan against a \$25 cash payment.[26]

In the following, we estimate individual-specific values of each type of grant, loan and hybrid loan that are proposed in the context of the experiment. Put into a public policy perspective, these estimates can be interpreted as uncovering the willingness-to-pay for a counterfactual expansion in higher education financial aid offers. In that sense, the experiment we use in the paper allows us to go beyond take-up rates by exploring the intensive margin of the demand for college financial aid.

We choose to pay specific attention to the willingness-to-pay for loans, which, unlike the willingness-to-pay for grants, are a natural measure of the tightness of borrowing constraints, as perceived by the students by the end of high school. While constrained students may attach a significant value to the opportunity of taking up a loan at the market rate, unconstrained students should only attach small or negligible values to those redundant opportunities.[27]

Since the experiment was designed to elicit the willingness-to-pay for higher education financial aid ahead of high school graduation, these estimates are best thought as *ex ante* measures that speak to the students' perceptions of, rather than realized, credit constraints. To the extent that these perceptions likely affect important decisions such as applying to and enrolling in college, as well as potentially also the rate of human capital accumulation while in high school, these *ex ante* measures are key to evaluating the importance of education financing barriers in the economy. Depending

[26] In order to examine the importance of allowing for heterogeneous time and risk preference parameters, we also have estimated a constrained specification of our model with homogeneous discount factors and relative risk aversion parameters. While the estimated discount factor and risk aversion parameters (equal to 0.78 and 0.67, respectively) are close to the mean of the corresponding empirical distributions from our preferred specification, the fit of this restricted model is much less satisfactory.

[27] In principle, the willingness-to-pay for loans may also capture the hassle cost associated with student loan application outside of the experiment. However, for the case of Federal student loans in Canada (Canada Student Loan Program, CSLP) the (online) application process requires filling one form only, and is overall fairly easy and quick. As a result, and in contrast with the Free Application for Federal Student Aid (FAFSA) process in the United States, the hassle cost seems unlikely to play a major role in this context. Irrespective of differences in hassle costs between CSLP loans and loans from the experiment, we cannot rule out that some students prefer the latter, potentially because they assume there are lower costs of obtaining it, or because they believe that the implications of not repaying it may not be as serious as with the CSLP loans.

on how accurate students' perceptions about their future borrowing constraints are, these *ex ante* measures may be more or less tightly associated with the *ex post* borrowing constraints.[28]

Specifically, let c_{iq}^m be the incremental level of consumption that makes individuals indifferent between current consumption $(c_i + c_{iq}^m)$ and the financial aid package at question q.[29] For such a value, we have

$$u(c_i, \theta_i) + \beta_i \psi_{iq} = u(c_i + c_{iq}^m, \theta_i) \tag{8}$$

After a few steps of algebra, we obtain:

$$c_{iq}^m = \begin{cases} \exp\left(\frac{1}{1-\theta_i} \log\left((1-\theta_i)\beta_i \psi_{iq} + c_i^{1-\theta_i}\right)\right) - c_i & \text{if } \theta_i \neq 1 \\ \exp\left(\beta_i \psi_{iq} + \log(c_i)\right) - c_i & \text{if } \theta_i = 1 \end{cases} \tag{9}$$

c_{iq}^m is the maximum consumption increase that one would be willing to trade in order to secure the financial aid package, or, in other words, the consumption equivalent of the added welfare of the financial aid. Throughout the rest of the paper, we use c_{iq}^m as our individual-specific measure of the willingness-to-pay for the financial aid package offered in question q.

7.3.2 Willingness-to-pay for financial aid

In Table 4 below, we summarize the distributions of the estimated willingness-to-pay for the three types of financial aid packages (loans, grants and hybrid loans) of sizes \$1,000 and \$2,000.[30] Our results indicate that the median high school student would be willing to forego a \$60.6 increase in current consumption to secure a \$1,000 loan at the market interest rate in the near future. The willingness-to-pay for a \$1,000 loan is highly heterogeneous across students, with an interdecile range equal to \$174. While a quarter of the students are willing to sacrifice more than \$116.7 for the option to take up a \$1,000 loan, students in the bottom quartile are only willing to sacrifice less than \$22.5.

As the experiment was conducted throughout the academic year, we can also examine how the willingness-to-pay for financial aid varies over time. In Table 5 below, we document how the perceptions of credit constraints vary over time by reporting the distributions of the estimated willingness-to-pay for loans of sizes \$1,000 and \$2,000, separately for the students who were interviewed between October and December, and those interviewed between January and March.

[28] Since loan conditions were similar to those of the Federal Canadian Student Loan Program, such (*ex post*) borrowing constraints may arise as some students exhaust their CSLP loan limits. In practice, over the year 2008-2009, a substantial share (37%) of the students who took up a federal loan reached the loan limit (Office of the Chief Actuary, 2009).

[29] In practice we set $c_i = c_i^3$ (background consumption level for the education financing questions).

[30] For the ease of exposition, we focus hereafter on financial aid packages of sizes \$1,000 and \$2,000. Estimation results for alternative amounts of financial aid are available from the authors upon request.

Table 4: The distribution of willingness-to-pay

	$1,000			$2,000		
	Loan	Grant	Hybrid	Loan	Grant	Hybrid
1st Dec.	4.6	68.3	71.3	5.1	160.9	140.9
1st Quart.	22.5	293.2	273.5	29.4	557.7	420.5
Median	60.6	449.8	409.0	72.2	792.4	578.5
3rd Quart.	116.7	638.8	566.0	129.7	1118.2	763.4
9th Dec.	178.6	795.0	682.1	182.3	1469.2	946.4

Notes: Amounts are in Canadian dollars.

Table 5: Willingness-to-pay for loans (time of the experiment)

	Oct-Dec		Jan-March	
	1,000 $	2,000 $	1,000 $	2,000 $
1st Dec.	8.2	8.2	4.7	3.8
1st Quart.	29.1	36.5	22.0	27.4
Median	72.4	85.1	60.3	74.9
3rd Quart.	129.8	137.0	124.2	139.0
9th Dec.	192.7	197.7	178.7	196.2

Notes: Amounts are in Canadian dollars.

We find that students who are interviewed in the first half of the school year are generally willing to give up larger amounts of cash payments for the option to take up a college loan at the prevailing market rate. For instance, the median willingness-to-pay for a $1,000 ($ 2,000) loan is equal to $72.4 ($85.1) among students who are interviewed between October and December, against $60.3 ($74.9) for those interviewed between January and March. Taken together, these results provide suggestive evidence that, as students learn about the financial aid opportunities during their senior year of high school, they tend to attach somewhat smaller values to the loans that are offered as part of the experiment. Importantly though, that the willingness-to-pay remains sizable among the students who are interviewed in the second half of the year indicates that learning about financial aid opportunities while in high school is not the only mechanism at play.

Put into a public policy perspective, a grant is equivalent to a tuition reduction, or a higher education subsidy. Not surprisingly, the value of a grant is typically much larger. The median student would be willing to trade in about $450 increase of their current consumption for the option to take up a $1,000 grant. Contrary to loans, only a small proportion of the population attaches

low values to grant availability. For instance, less than 10% of the students would sacrifice less than \$68 for a \$1,000 grant, while more than half of them would do so for a \$1,000 loan. More generally, for both amounts of financial aid (\$1,000 and \$2,000), the distribution of the willingness-to-pay for a grant stochastically dominates that of a loan.[31]

Turning to the hybrid packages, adding a loan to a grant generally has a small negative impact on the value of the package. This pattern holds true in most parts of the distributions of the willingness-to-pay. These results likely reflect the fact that, in practice, students can't receive the grant without the loan component. It follows that those students who end up taking up the hybrid package upon college entry also have to pay the loan back with interest. As a result, how much students are willing to trade for a hybrid package versus a single grant depends on the incremental value of a loan, as well as on the interest cost associated with it. Our results indicate that, in this context, the latter effect dominates. At any rate, this provides additional evidence that our model fits the descriptive patterns previously reported in Section 3 (Figures 1 and 2).

Finally, for loans, grants as well as hybrid loans, the value of the package generally increases with the size of the financial aid. Specifically, the distribution of the willingness-to-pay for a \$2,000 loan (grant) stochastically dominates that of a \$1,000 loan (grant), while a (\$2,000 loan, \$2,000 grant) hybrid loan also dominates that of a (\$1,000 loan, \$1,000 grant) hybrid loan. In particular, the results for loans provide suggestive evidence that, at least for a subset of the students in the sample, getting access to a \$1,000 loan is not enough to remove higher education credit constraints.

One can quantify the magnitude of, and heterogeneity in credit market imperfections by converting the willingness-to-pay for any given loan into the interest rate wedge that students would be willing to pay on top of the prevailing ($r_0 = 5.7\%$) market interest rate in order to secure the option to take up that loan. Specifically, for a given loan of size l and willingness-to-pay c_m, we define the effective interest rate as the interest rate, denoted by $r_1(l, c_m)$, such that:

$$(1 + r_1(l, c_m))\,(l - c_m) = (1 + r_0)l \tag{10}$$

The effective interest rate is defined as the interest rate associated with a total repayment $(1+r_0)l$ (total repayment for a loan l at the prevailing market rate), and a principal $l - c_m$ (principal of the loan l net of the willingness-to-pay for that loan).[32] Note that $r_1(l, c_m)$, which is an increasing function of the willingness-to-pay c_m, is larger than the prevailing rate r_0 for any positive c_m.

We report in Table 6 below the distribution of the interest rate wedges $(r_1(l, c_m) - r_0)$ associated with a \$1,000 loan and a \$2,000 loan:

[31] Both loans and grants might crowd-out parental transfers. Such a response would be captured in our estimates of the future value of, and willingness-to-pay for financial aid, provided that it is anticipated by the students by the end of high school.

[32] This definition abstracts from the fact that c_m is measured a few months before the loan is disbursed. Taking this lag into account when computing the effective interest rate - which would require additional assumptions in terms of timing and savings interest rate - would result in marginally larger interest rate wedges.

Table 6: The distribution of interest rate wedges (in percentage points)

Quantiles	$1,000 loan	$2,000 loan
1st Dec.	0.49	0.27
1st Quart.	2.43	1.58
Median	6.82	3.96
3rd Quart.	14	7.33
9th Dec.	23	10.6

Note: Interest rate wedges are computed with respect to the prevailing market interest rate of 5.7%.

These results indicate that most of the students in our sample are willing to pay a sizable interest rate premium above the prevailing market rate for the option to take up a loan. The median interest rate premium students would be willing to pay is large and equal to 6.82 (3.96) percentage points for a $1,000 ($2,000) loan, respectively.[33] As expected given the results on the willingness-to-pay for loans documented earlier, the interest rate premia also exhibit much heterogeneity across students. For the case of $1,000 loans, the interest wedge ranges from 0.49 to 23 percentage points for the first and the last decile, respectively. Overall, these results point to the existence of credit constraints, in the form of frictions in the market for college loans, which affect a substantial share of high school students in Canada. Our estimates further show that a non-negligible fraction of high school students attach pretty large values to the option to take up a college loan. From a policy standpoint, our findings indicate that expanding higher education financial aid may in fact be socially desirable in this context, in spite of the Canadian higher education system being already heavily subsidized.[34]

Finally, it is worth noting that, since some of the students in the sample expect not to enroll in college, our results effectively provide a lower bound on the willingness-to-pay for financial aid conditional on expecting to go to college. We address this issue by repeating our analysis on the subsample of students who declare that they expect (or aspire) to go to college.[35] The results are reported in Table A16 in Appendix H.1. As expected, willingness-to-pay for financial aid tends to be larger for the students who expect (or aspire) to attend college. Expressed in terms of median interest rate wedges for a $1,000 loan, restricting to the group of students who expect (aspire) to go

[33] It follows from Equation (10) that interest wedges are increasing with the willingness-to-pay c_m, but decreasing with the loan amount l. That the distribution of the interest wedges for a $1,000 loan dominates the distribution associated with a $2,000 loan is evidence that, in this context, the latter effect dominates.

[34] In Appendix G, we report the distribution of the willingness-to-pay separately for Ontario and Québec. Our estimates are consistent with the institutional differences across these two provinces. In particular, the distribution of the willingness-to-pay for a $1,000 loan in Ontario dominates almost everywhere the distribution in Québec, consistent with higher levels of (net) tuition fees in Ontario than in Québec.

[35] In Table A20 in Appendix I we examine how these measures of higher education attendance expectation and aspiration vary by background characteristics, numeracy test score as well as risk and time preference parameters.

to college results in an increase from 6.8 to 8.2 (8.1) percentage points.

7.3.3 Explaining the willingness-to-pay

We now turn to the marginal effects of parental background variables and other observed attributes and preferences on the willingness-to-pay for financial aid. Table 7 below reports the average marginal effects for two types of financial aid packages, namely a $1,000 loan and a $1,000 grant.[36]

We first examine the effects of parental income. There is a large empirical literature on the relationship between parental income and schooling attainment (Heckman and Mosso, 2014). Economists have long debated on the magnitude of the causal effects of parental income on educational outcomes, in particular college enrollment. Identifying those effects is a complicated task as family income is also likely correlated with individual abilities, as well as preferences for schooling. The experiment used in this paper allows us to go beyond evaluating the impact of parental income on schooling attainment as we can directly quantify how the willingness-to-pay for education financial aid, a measure that increases with the anticipated intensity of credit constraints, vary with family income. Doing so is also an important step towards evaluating the effectiveness of publicly provided financial aid policies that are meant to equalize opportunities across income groups. The results reported in Table 7 illustrate the differences in willingness-to-pay across income classes, using as a reference those who earn $20,000 or less. The results indicate that the willingness-to-pay for a $1,000 loan is non-linear and non-monotonic as the highest willingness-to-pay is found for the $40,000-$60,000 income group. However, the marginal effects across income groups are small. For instance, on average, young individuals raised in families earning $40,000-$60,000 would be willing to pay $1.2 more for the option to take up a $1,000 loan, relative to the lowest income reference group. A similar pattern holds for grants. The highest willingness-to-pay for a $1,000 grant is found for students from the $60,000-$80,000 income group, who would pay $1.5 more than those from the reference group. The marginal effects of parental education on the willingness-to-pay for a $1,000 loan and a $1,000 grant are also generally very small.

The co-existence of positive and sizable values attached to loans and grants, documented in Section 7.3.2 (Table 4), with the quasi-independence of the willingness-to-pay with respect to parental income indicates two important features of the Canadian higher education financing system. First, the median Canadian high school student is not satiated with financial aid opportunities, and the expected marginal utility of financial aid is non-negligible. Second, it does appear that the higher education public policies in place in Canada are successful in equalizing the marginal utility of financial aid opportunities across various income groups, as differences in family income have no impact on the value attached to a counterfactual expansion in higher education financial aid

[36] The fact that the magnitudes of the marginal effects tend to be larger for a $1,000 grant than for a $1,000 loan partly reflects the difference in scale across both outcomes, with the standard deviation of the willingness-to-pay for a $1,000 grant being equal to $223, against $77 for a $1,000 loan.

Table 7: Explaining the willingness-to-pay (Average Marginal Effects)

	$c_m^i(\ell_{1000})$	$c_m^i(g_{1000})$
θ	0.23	1.60***
	(0.17)	(0.13)
β	3.90***	14.53***
	(0.26)	(0.30)
Numeracy	0.70	0.50***
	(2.63)	(0.08)
20-40K	-0.56	-0.44
	(0.15)	(0.21)
40-60K	1.23***	-0.37
	(0.16)	(0.08)
60-80K	-0.34	1.49***
	(0.29)	(0.50)
80-100K	-0.39	-0.02
	(0.08)	(0.05)
+100K	0.06	-0.39
	(0.05)	(0.09)
High-school	-0.38	0.94***
	(0.08)	(0.15)
V/College	0.13	1.01**
	(0.65)	(0.49)
College	-0.21	0.39*
	(0.32)	(0.22)
Quebec	-0.24	-0.16
	(0.06)	(0.03)
Manitoba	0.45***	1.12***
	(0.10)	(0.14)
Saskatchewan	0.90***	-0.59
	(0.18)	(0.08)

Notes: Notes: (i) The entries in this table are computed as the sample averages of the marginal effects evaluated at the observed values of the vector of covariates (Average Marginal Effects), (ii) marginal effects in standard deviation units for all continuous characteristics (θ, β, and Numeracy), (iii) standard errors reported in parenthesis, (iv) $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

opportunities. The lack of a significant relationship between parental income and willingness-to-pay for financial aid also appears to be consistent with prior evidence from the literature of a weak family income - post-secondary attendance gradient in Canada (Belley et al., 2014).

Ultimately, individual differences in the willingness-to-pay for financial aid are mostly explained by deep (time and risk) preference parameters. We now discuss their marginal effects. The effect of the discount factor is easily predictable as the structure of the experiment implies that the benefit of financial aid can only be experienced in the future while the cash payment is practically immediate. Indeed, as shown in Table 7, the marginal effect of the discount factor on the willingness-to-pay for grants and loans are both positive and significant. All else equal, increasing the discount factor by one standard deviation leads to a $3.9 ($14.5) increase in the willingness-to-pay for a $1,000 loan ($1,000 grant). These marginal effects are in particular substantially larger than any of the effects associated with a $20,000 income differential.

The effect of risk aversion on the willingness-to-pay for grants or loans is more intricate because risk aversion affects not only the value of accepting a cash payment but also the value of financial aid, with the effect of risk aversion on the value of financial aid being ambiguous. Indeed, while high-school students who are more risk averse are also those who would benefit more from the future consumption smoothing opportunity provided by financial aid, individuals who accept financial aid can only exercise the option to take up a loan or a grant if they end up enrolling in higher education. Those who are more risk averse and who are also not certain of entering higher education, or are not sure when they are likely to do so, will be particularly sensitive to the latter feature. As a consequence, it is in theory not possible to sign unambiguously the effect of risk aversion on the willingness-to-pay for financial aid. Because both individual-specific enrollment uncertainty and future parental transfers are unobserved in our context, our model is not capable of separating those two channels, and, while interesting, any discussion about their relative importance would lie beyond the scope of the paper. At this stage, we only note that the marginal effect of a one standard deviation increase in risk aversion on the willingness-to-pay for a $1,000 grant and a $1,000 loan are both positive, but are 9 to 17 times smaller than the marginal effects of discount factors.

7.3.4 Robustness checks

We conclude this section by documenting the robustness of our main findings to alternative estimation samples and specifications.

In Table A14 in Appendix H.1, we show that the distribution of the estimated willingness-to-pay for financial aid is robust to (i) the exclusion from the estimation sample of respondents who have very low discount rates ($\beta_i \leq 0.1$), as well as to (ii) the exclusion of respondents with risk aversion parameters below 0.19 and over 3.5, which correspond to the 2.5 and 97.5 percentiles of the distribution, respectively. In this table we also show that our results remain largely unchanged after

(iii) excluding the set of financial aid questions that are associated with the highest cash payment ($700), for which the maintained assumption that students do not smooth their consumption gain while in high school is, at least *a priori*, more restrictive.

Next, we report in Table A15 the distribution of the willingness-to-pay for financial aid obtained after excluding from the estimation sample the financial aid questions that involve grants. By doing so, we infer the willingness-to-pay for loans using loan-related questions only. Estimation results are again very similar to our baseline results.

In Tables A17 and A18 in Appendix H.2, we report the distribution of the willingness-to-pay for financial aid ($1,000 and $2,000, respectively), for our baseline specification ("Benchmark") and three alternative specifications of the future value component (ψ_{iq}). The second specification ("Parsimonious") corresponds to a more parsimonious specification where all variables enter linearly the future value component. The third specification ("Expanded") corresponds to a more flexible specification where all of the variables are interacted. Finally, in the fourth and last specification ("Shock"), we include the idiosyncratic shocks ε_{iq} additively in the future value component ψ_{iq}. For all three alternative specifications, estimation results remain again similar to the benchmark specification, both qualitatively and in many cases quantitatively.

Finally, in Table A19 in Appendix H.3, we report the distribution of the willingness-to-pay for financial aid obtained when we allow background consumption to also depend on individual skills, as measured by the numeracy score. We also consider an expanded specification where all covariates from the baseline specification of background consumption are interacted. Our results remain robust to these two alternative specifications.

8 Conclusion

In this paper, we estimate the distribution of the willingness-to-pay for higher education financial aid using data from a field experiment conducted in Canada where high school students had to choose between immediate cash payments and various types of higher education financing packages. Our model of financial aid acceptance decisions is based on an explicit trade-off between the increase in current utility following an immediate cash payment and the expected future gain associated with a specific financial aid package. As the experiment also allows us to estimate the distributions of individual risk aversion and discount factors, we are able to uncover how preference parameters for time and risk affect the willingness-to-pay for the various types of financial aids.

We find that the majority of students attach a sizable value to accessing student loans. Consistent with the existence of significant frictions in the market for college loans, we find that the median high school student in our sample would be willing to pay a large 6.8 percentage points interest rate premium on top of the prevailing market rate to secure a $1,000 loan. Taken together, these

findings point to the existence of credit constraints affecting a sizable share of the population of high school students.

Data availability has confined us to the analysis of financing decisions prior to actual college enrollment. It would be interesting to combine this educational financing experiment with observational data on subsequent outcomes to gain additional insights into the effects of improving higher education financing opportunities on educational as well as future labor market outcomes. This is left for future research.

References

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

Andreoni, J. and C. Sprenger (2012). Estimating time preferences from convex budgets. *American Economic Review 102*(7), 3333–3356.

Arcidiacono, P., V. J. Hotz, and S. Kang (2012). Modeling college major choices using elicited measures of expectations and counterfactuals. *Journal of Econometrics 166*(1), 3–16.

Arcidiacono, P., H. Sieg, and F. Sloan (2007). Living rationally under the volcano? an empirical analysis of heavy drinking and smoking. *International Economic Review 48*, 37–65.

Belley, P., M. Frenette, and L. Lochner (2014). Post-secondary attendance by parental income in the U.S. and Canada: do financial aid policies explain the differences? *Canadian Journal of Economics 47*(2), 664–696.

Belley, P. and L. Lochner (2007). The changing role of family income and ability in determining educational achievement. *Journal of Human Capital 1*(1), 37–89.

Belzil, C. and J. Hansen (2004). *Earnings Dispersion, Risk Aversion and Education*, Volume 33 of *Research in Labor Economics*, pp. 335–358. Greenwich, Conn.: JAI.

Brodaty, T., R. J. Gary-Bobo, and A. Prieto (2014). Do risk aversion and wages explain educational choices? *Journal of Public Economics 117*, 125–148.

Brown, M., J. K. Scholz, and A. Seshadri (2012). A new test of borrowing constraints for education. *Review of Economic Studies 79*(2), 511–538.

Cameron, S. V. and J. J. Heckman (1998). Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of american males. *Journal of Political Economy 106*(2), 262–333.

Cameron, S. V. and C. Taber (2004). Estimation of educational borrowing constraints using returns to schooling. *Journal of Political Economy 112*(1), 132–182.

Card, D. (1995). *Earnings, schooling, and ability revisited*, Volume 14 of *Research in Labor Economics*, pp. 23–28. Greenwich, Conn.: JAI.

Carneiro, P., K. T. Hansen, and J. J. Heckman (2003). Estimating distributions of treatment effects with an application to the returns to schooling and measurement of the effects of uncertainty on college choice. *International Economic Review 44*(2), 361–422.

Carneiro, P. and J. J. Heckman (2002). The evidence on credit constraints in post–secondary schooling. *Economic Journal 112*(482), 705–734.

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

Delavande, A. and B. Zafar (2019). University Choice: The Role of Expected Earnings, Non-pecuniary Outcomes and Financial Constraints. *Journal of Political Economy 127*(5), 2343–2393.

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time discounting and time preference: a critical review. *Journal of Economic Literature 40*(2), 351–401.

Geweke, J. and M. Keane (2000). Bayesian inference for dynamic discrete choice models without the need for dynamic programming. In R. Mariano, T. Schuermann, and M. Weeks (Eds.), *Simulation-based inference in econometrics: methods and applications*, pp. 100–131. Cambridge: Cambridge University Press.

Hai, R. and J. J. Heckman (2017). Inequality in human capital and endogenous credit constraints. *Review of Economic Dynamics 25*, 4–36.

Hansen, J. and X. Liu (2013). A structural model of educational attainment in canada. IZA Discussion Paper 7237.

Heckman, J. J. and S. Mosso (2014). The economics of human development and social mobility. *Annual Review of Economics 6*, 689–733.

Holt, C. A. and S. K. Laury (2002). Risk aversion and incentive effects. *American Economic Review 92*(5), 1644–1655.

Hotz, V. J., E. Wiemers, J. Rasmussen, and K. Maxwell Koegel (2018). The Role of Parental Wealth and Income in Financing Children's College Attendance and its Consequences. NBER Working Paper 25144, National Bureau of Economic Research, Inc.

Humphries, J. E. (2018). The causes and consequences of self-employment over the life cycle. Working paper, Yale University.

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

Johnson, M. T. (2013). Borrowing constraints, college enrollment, and delayed entry. *Journal of Labor Economics 31*(4), 669–725.

Keane, M. P. and K. I. Wolpin (1997). The career decisions of young men. *Journal of Political Economy 105*(3), 473–522.

Keane, M. P. and K. I. Wolpin (2001). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review 42*(4), 1051–1103.

Lang, K. (1993). Ability bias, discount rate bias, and the return to education. Manuscript, Boston University, Department of Economics.

Lochner, L., T. Stinebrickner, and U. Suleymanoglu (2018). Parental Support, Savings and Student Loan Repayment. NBER Working Paper 24863, National Bureau of Economic Research.

Lochner, L. J. and A. Monge-Naranjo (2011). The nature of credit constraints and human capital. *American Economic Review 101*(6), 2487–2529.

Lochner, L. J. and A. Monge-Naranjo (2012). Credit constraints in education. *Annual Review of Economics 4*, 225–256.

Lochner, L. J. and A. Monge-Naranjo (2016). Student loans and repayment: theory, evidence, and policy. Volume 5 of *Handbook of the Economics of Education*, Chapter 8, pp. 397–478. Amsterdam: North Holland.

Magnac, T. and E. Maurin (2008). Partial identification in monotone binary models: Discrete regressors and interval data. *Review of Economic Studies 75*(3), 835–864.

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

Murray, S., Y. Clermont, and M. Binkley (2005). The international adult literacy and skills survey. Technical report, Statistics Canada.

Neyman, J. and E. L. Scott (1948). Consistent estimates based on partially consistent observations. *Econometrica 16*(1), 1–32.

Office of the Chief Actuary (2009). Actuarial report on the canada student loans program as at 31 july 2009. Technical report.

Pavan, R. (2011). Career choice and wage growth. *Journal of Labor Economics 29*(3), 549–87.

Rust, J. and C. Phelan (1997). How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica 65*(4), 781–832.

Sauer, R. M. (2004). Educational financing and lifetime earnings. *Review of Economic Studies 71*(4), 1189–1216.

Stinebrickner, R. and T. R. Stinebrickner (2008). The effect of credit constraints on the college drop-out decision: A direct approach using a new panel study. *American Economic Review 98*(5), 2163–84.

Stinebrickner, R. and T. R. Stinebrickner (2014). A major in science? initial beliefs and final outcomes for college major and dropout. *Review of Economic Studies 81*(1), 426–472.

Wiswall, M. and B. Zafar (2015a). Determinants of college major choice: identification using an information experiment. *Review of Economic Studies 82*(2), 791–824.

Wiswall, M. and B. Zafar (2015b). How do college students respond to public information about earnings? *Journal of Human Capital 9*(2), 117–169.

Zafar, B. (2013). College major choice and the gender gap. *Journal of Human Resources 48*, 545–595.

Appendix

A Higher education financial aid in Canada

The main forms of student financial aid in Canada are student loans and grants. We briefly review these two types of financial aid in the next two subsections.

A.1 Loans

A key part of higher education financing in Canada comes from the Canada Student Loan Program (CSLP), which is funded by the federal government. The main eligibility criteria are the following:[37]

1. Be a Canadian citizen, a permanent resident of Canada or designated as a protected person.
2. Be a permanent resident of a province or territory that issues Canada Student Loans.
3. Demonstrate financial need.
4. (For a full-time student): be enrolled in at least 60 percent of a full course load (40 percent for students with permanent disabilities).
5. (For a part-time student): be enrolled in 20-59 percent of a full course load (for student with permanent disabilities and studying between 40-59 percent of a full course load, one can choose to be considered a student in full- or part-time studies).
6. Be enrolled in a degree, diploma or certificate program offered by a designated post-secondary school that runs for at least 12 weeks within a 15-week period.
7. Pass a credit check if 22 or older and are applying for the first time.
8. Not have exhausted one's maximum lifetime limit for financial assistance (including interest-free status).[38]

The application process itself requires filling one (online) form only, and is overall fairly easy and quick. The documents that need to be provided as part of the application process include a valid photo identification issued in Canada, a document justifying the Social Insurance Number, and a signed Student Financial Assistance Agreement.

[37] Source: https://www.canada.ca/en/employment-social-development/services/student-financial-aid/student-loan/student-loans/eligibility.html.

[38] In practice, lifetime limits for Canada Student Loans vary across candidates. Namely: (i) full-time students who received loans on or after August 1, 1995, are eligible to receive student financial assistance for no more than 340 weeks; (ii) full-time students enrolled in doctoral studies are eligible to receive student financial assistance for no more than 400 weeks; and (iii) students with either a permanent disability or who received Canada Student Loans before August 1, 1995, are eligible to receive student financial assistance for no more than 520 weeks.

Students who are eligible and take up a loan as part of the CSLP have to start paying back the loan six months after graduation or college exit.

Additional types of (non-Government) issued loans are also available to students. In particular, financial institutions typically offer regular consumer loans for school, lines of credit as well as emergency loans. The differences between non-government and government student loans mainly lie in the requirement of guarantees, interests, the time settings of paying back the loan, and the forms of repayment assistance.

A.2 Grants

The CSLP also offers grants, that are available to students from most provinces and territories except the Northwest Territories, Nunavut and Quebec which have their own student financial assistance programs.

The CSLP offers these different types of grants:

1. Grant for students from low-income families: the core eligibility criterion is to be from a low-income family, as defined by the CSLP. Students can receive up to $375 per month, for each year of undergraduate studies.

2. Grant for students from middle-income families: the core eligibility criterion is to be from a middle-income family, as defined by the CSLP. Students can receive up to $150 per month, for each year of undergraduate studies.

3. Grant for full-time students with dependents: the core eligibility criteria are to be from a low-income family as defined by the CSLP, and having a dependent who will be under 12 years of age at the time of enrollment (or a dependent 12 years of age or older with a permanent disability). Students in this category can receive up to $200 per month for every dependent child for each year of undergraduate and graduate studies, in addition to any amount received under the Canada Student Grant for students from low-income families.

4. Grant for part-time studies: the core eligibility criteria include being enrolled in a part-time degree, diploma, or certificate program (at least 12 weeks long within a period of 15 weeks in a row) at a designated post-secondary institution, and the student and his/her spouse or common-law partner have a low family income as defined by the CSLP. The amounts can reach $1,800 per year of undergraduate and graduate studies.

5. Grant for part-time students with dependents: the core eligibility criteria include: (i) being enrolled in a part-time degree, diploma, or certificate program (at least 12 weeks long within a period of 15 weeks in a row) at a designated post-secondary institution; (ii) having a dependent

who will be under 12 years of age at the beginning of the study period (or a dependent with a permanent disability); and (iii) the student and his/her spouse or common-law partner have a low family income as defined by the CSLP. This grant is in addition to any grant received under the Canada Student Grant for part-time studies, and the amount one could receive is a maximum of $40 per week (for students with at most 2 dependents), or up to a maximum of $1,920 per year of undergraduate and graduate studies (for students with 3 dependents or more).

6. Grants for students with permanent disabilities: qualification for this grant requires meeting the criteria for students with permanent disabilities, including providing proof of disability. The maximum amount is $2,000 per school year (August 1 to July 31) of undergraduate and graduate studies.

7. Grant for services and equipment for students with permanent disabilities: eligibility requires providing written confirmation that one is in need of exceptional education-related services or equipment from a person qualified to determine such need, along with written confirmation of the exact cost of the equipment and services. The maximum amount is $8,000 per school year.

Provincial and territorial grants and bursaries are also available to the students. Finally, another source of financial aid is the Canada Education Savings Grant (CESG). The CESG corresponds to money the Government adds to one's child's Registered Education Savings Plan to help their savings grow. After high school, students can withdraw the money to help pay for either full-time or part-time studies in an apprenticeship program, at a CEGEP, at a trade school, at a college, or at a university. The basic CESG provides 20 cents on every dollar contributed, up to a maximum of $500 on an annual contribution of $2,500, and is available up until the end of the calendar year in which the child turns 17. Depending on the child's primary caregiver's net family income, he or she may also be eligible to receive the Additional Canada Education Savings Grant. This grant adds an additional 10% or 20% to the first $500 put into the RESP each year.

B The Millenium Foundation Field Experiment on Education Financing

In this section, we provide additional details on the structure of the field experiment. An extensive discussion of the experiment and the sample is available in Johnson and Montmarquette (2015).

All subjects were presented with the full set of decisions and are paid for one randomly selected decision 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 can be split into three groups. First, students must answer a set of questions aimed at measuring their

rate of time preference. Table A1 illustrates the experiment where individuals are offered a choice between two payments of different values to be made at different points in time. This approach,

Table A1: Discount rates questions

Choices 1		Choices 2		
Payment 1 day	Payment 1 month	Payment 1 day	Payment 1 year	Annual Interest(%)
$ 75	$ 75.31	$ 75	$ 78.75	5
$ 75	$ 75.63	$ 75	$ 82.5	10
$ 75	$ 76.25	$ 75	$ 90.00	20
$ 75	$ 78.13	$ 75	$ 112.5	50
$ 75	$ 81.25	$ 75	$ 150.0	100
$ 75	$ 87.5	$ 75	$ 225.0	200
Choices 3		Choices 4		
Payment 1 week	Payment 1 month + 1 week	Payment 1 week	Payment 1 year + 1 week	Annual Interest(%)
$ 75	$ 75.31	$ 75	$ 78.75	5
$ 75	$ 75.63	$ 75	$ 82.5	10
$ 75	$ 76.25	$ 75	$ 90.00	20
$ 75	$ 78.13	$ 75	$ 112.5	50
$ 75	$ 81.25	$ 75	$ 150.0	100
$ 75	$ 87.5	$ 75	$ 225.0	200

originally developed by Coller and Williams (1999), allows to evaluate each subject's willingness to forgo present consumption for future consumption, and provides a measure of the discount rate.

A second set of questions relate to the measurement of risk attitudes. Students are presented with a sequence of binary choices between two lotteries in which risk is objectively stated. Table A2 presents the two strategies used for eliciting risk aversion. Both strategies consist of choosing between a lottery with average payoff and another one with extreme payoff, and identify the cutoff point where an agent switches from the average to the extreme lottery. The major difference between the two strategies lies in the fact that while the first one pins down a cut-off probability, the second identifies a cut-off payoff. These approaches pioneered by Holt and Laury (2002) are standard in the experimental literature to measure the degree of risk aversion.

The third group of questions, which constitutes the most original aspect of the field experiment, is a sequence of choices between a cash payment to be paid within one week from the day the experiment was carried, and the option to use a specific financial aid package covering educational expenses. The financial aid package is to be paid conditional on enrolling in a full-time program at

Table A2: Risk aversion lotteries

Panel 1: Moving probability Lotteries (30 questions)

			Choices 1		Choices 2		Choices 3	
			L= 32	L=2	L=24	L=1.5	L=40	L=2.5
			H=40	H=77	H=30	H=57.75	H=50	H=96.25
	Prob L	Prob 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

Panel 2: Fixed probability Lotteries (Prob L= Prob H = 0.5, 25 questions)

L	H	$\mathbb{E}V$	L	H	$\mathbb{E}V$	L	H	$\mathbb{E}V$	L	H	$\mathbb{E}V$
Options A			Options B			Options A			Options B		
48	48	48	40	64	52	18	90	54	8	104	56
40	64	52	32	80	56	8	104	56	0	112	56
32	80	56	24	96	60	42	42	42	36	60	48
24	96	60	16	112	64	36	60	48	30	78	54
16	112	64	8	120	64	30	78	54	24	96	60
48	48	48	42	66	54	24	96	60	18	114	66
42	66	54	36	84	60	18	114	66	10	122	66
36	84	60	30	102	66	54	54	54	44	68	56
30	102	66	24	120	72	44	68	56	34	82	58
24	120	72	16	128	72	34	82	58	24	96	60
48	48	48	38	62	50	24	96	60	14	110	62
38	62	50	28	76	52	14	110	62	6	118	62
28	76	52	18	90	54						

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.

any higher education institution in the country (within 2 years).

In total, we consider three different types of subsidies: grants, loans, and hybrid loans which incorporate both a loan and a grant component. We use a total of 17 financial decisions, with 5 choices with a single loan offer, 7 choices with a single grant offer, and 5 hybrid offers. These decisions are summarized in Table 1.

In monetary terms, cash alternatives varied from $25 to $700, while grants and loans varied from $400 to $4,000. The variations in cash, and in the nature and the size of financial aid packages have a number of advantages. For instance, for a given cash payment offered and manipulating the financial parameters, we can uncover the relative values of a grant and a loan. Suppose instead that the financial aid package is fixed, we can also reveal the willingness to pay for a specific package by manipulating the cash payment.

C Descriptive characteristics of the sample

The sample was recruited to generate meaningful comparisons by population group, gender, and low-, medium- or high-income status. The original project design called for a minimum sample size of 1,000 urban respondents with the goal of 200 participants per group of interest, with a total sample of 1,248 individuals. Table A3 summarizes the share of participants in several groups of interest.[39]

[39] Parental education is defined as the highest level of education of the parent who responded to the survey. Parental income is defined as the total income declared by the respondent, before tax deductions and measured in 2007, of all family members living in the household.

Table A3: Descriptive statistics of the sample

Parental	Drop-out	8%
Education	High-school	25%
	Vocational College	7%
	College	61%
Parental	0-20K	5%
Income	20-40K	13%
	40-60K	23%
	60-80K	19%
	80-100K	15%
	+100K	24%
Location	Urban	80%
	Rural	20%
Province	Quebec	30%
	Ontario	29%
	Manitoba	28%
	Saskatchewan	13%
Citizen	Native	94%
	Immigrant	6%
Gender	Male	46%
	Female	54%

Source: SRDC-CIRANO Field Experiment on Education Financing.

Take-Up Rates Across Demographic Groups Table A4 describes financial aid take-up by some basic socio-demographic characteristics. A couple of remarks are in order. First, female students are more likely to choose a financial aid package compared to their male counterparts. While the differences by gender are small for some decisions (e.g. 700$ cash versus 2,000$ loan), female students are 18 points more likely to take-up a grant of 1,000$ against 700$ of cash. We observe similar patterns for immigrants and students from rural locations.

Table A4: Financial Aid Take-up (socio-demographic characteristics)

Choices	All	Gender		Citizenship		Location	
		Male	Female	Immigrant	Native	Urban	Rural
c_{25} VS ℓ_{2000}	0.46	0.41	0.50	0.53	0.45	0.44	0.55
c_{300} VS ℓ_{2000}	0.17	0.13	0.21	0.16	0.17	0.15	0.26
c_{700} VS ℓ_{2000}	0.05	0.05	0.05	0.06	0.05	0.05	0.05
c_{300} VS ℓ_{1000}	0.11	0.08	0.13	0.13	0.11	0.10	0.15
c_{300} VS ℓ_{4000}	0.28	0.25	0.31	0.33	0.28	0.26	0.36
c_{25} VS $\ell g_{1000} + g_{1000}$	0.83	0.79	0.87	0.87	0.83	0.83	0.86
c_{300} VS $\ell g_{1000} + g_{1000}$	0.64	0.56	0.70	0.78	0.63	0.64	0.61
c_{700} VS $\ell g_{1000} + g_{1000}$	0.39	0.32	0.45	0.42	0.39	0.39	0.37
c_{300} VS $\ell g_{400} + g_{400}$	0.29	0.21	0.35	0.38	0.28	0.28	0.31
c_{300} VS $\ell g_{2000} + g_{2000}$	0.73	0.66	0.79	0.82	0.72	0.73	0.73
c_{25} VS g_{1000}	0.89	0.85	0.92	0.90	0.89	0.88	0.90
c_{100} VS g_{1000}	0.83	0.78	0.87	0.90	0.82	0.82	0.84
c_{300} VS g_{1000}	0.69	0.62	0.75	0.81	0.68	0.69	0.67
c_{700} VS g_{1000}	0.41	0.32	0.50	0.38	0.41	0.41	0.44
c_{300} VS g_{500}	0.38	0.29	0.47	0.43	0.38	0.38	0.42
c_{300} VS g_{2000}	0.76	0.71	0.81	0.86	0.76	0.76	0.76
c_{300} VS g_{4000}	0.84	0.80	0.87	0.87	0.83	0.83	0.84

Notes: Financial aid take-up rates by basic socio-demographic attributes.

We now examine the importance of geography in the decision to accept financial aid packages. As reported before there are substantial spatial heterogeneity in income and tuition, and as such we expect students from different regions to react differently to financial packages. Students from Ontario are the most likely to accept financial aid packages, which is consistent with the high level of tuition. On the contrary, students from Saskatchewan are the least likely to opt for a financial aid packages, a choice that may be rationalized by the relative high earnings - low tuition situation of the province.

Regarding the effect of parental education, we find that students from college educated parents are more likely to choose a financial aid package, which may reflect their high probability of attending higher education. We should note though that there is not a clear gradient between parental education and financial aid take-up.

Table A5: Financial Aid Take-up (geography and income)

Choices	All	Provinces				Parental Education			
		Que--bec	Onta--rio	Mani--toba	Saskat--chewan	Drop--Out	High--School	Voc--College	Col--lege
c_{25} VS ℓ_{2000}	0.46	0.50	0.49	0.42	0.38	0.48	0.48	0.41	0.45
c_{300} VS ℓ_{2000}	0.17	0.17	0.21	0.16	0.12	0.17	0.18	0.10	0.18
c_{700} VS ℓ_{2000}	0.05	0.05	0.06	0.06	0.02	0.05	0.06	0.03	0.05
c_{300} VS ℓ_{1000}	0.11	0.11	0.13	0.11	0.08	0.10	0.12	0.08	0.11
c_{300} VS ℓ_{4000}	0.28	0.28	0.33	0.27	0.23	0.28	0.28	0.27	0.29
c_{25} VS $\ell g_{1000} + g_{1000}$	0.83	0.90	0.88	0.80	0.65	0.80	0.80	0.77	0.86
c_{300} VS $\ell g_{1000} + g_{1000}$	0.64	0.65	0.70	0.65	0.45	0.54	0.61	0.58	0.67
c_{700} VS $\ell g_{1000} + g_{1000}$	0.39	0.37	0.44	0.42	0.27	0.29	0.38	0.28	0.42
c_{300} VS $\ell g_{400} + g_{400}$	0.29	0.29	0.32	0.32	0.16	0.24	0.27	0.20	0.31
c_{300} VS $\ell g_{2000} + g_{2000}$	0.73	0.74	0.79	0.72	0.56	0.61	0.71	0.64	0.76
c_{25} VS g_{1000}	0.89	0.93	0.93	0.87	0.73	0.86	0.86	0.81	0.91
c_{100} VS g_{1000}	0.83	0.88	0.87	0.80	0.66	0.78	0.80	0.76	0.85
c_{300} VS g_{1000}	0.69	0.70	0.76	0.69	0.51	0.54	0.67	0.62	0.72
c_{700} VS g_{1000}	0.41	0.39	0.46	0.45	0.30	0.32	0.41	0.33	0.43
c_{300} VS g_{500}	0.38	0.37	0.43	0.43	0.23	0.25	0.36	0.35	0.42
c_{300} VS g_{2000}	0.76	0.78	0.83	0.74	0.62	0.65	0.73	0.67	0.80
c_{300} VS g_{4000}	0.84	0.84	0.90	0.82	0.70	0.77	0.82	0.77	0.86

Notes: Financial aid take-up by parental location and income.

Take-Up Rates and Expectations/Aspirations We now analyze the effect of expectations and aspirations about higher education outcomes of financial aid take-up decisions. We use information from the students survey, which was administrated prior to the financial decision part. Specifically, the survey asks: "As things stand now, what is the highest level of education you think you will get?". Students are also asked an additional question: "What is the highest level of education you would like to get?". We use these questions to elicit expectations and aspirations about higher education.

As expected, Table A6 below shows larger financial aid uptake for individuals who expect (or aspire to) pursue higher education.

Table A6: Financial Aid Take-up (expectation and aspiration for college attendance)

Choices	All	Expectation		Aspiration	
		No	Yes	No	Yes
c_{25} VS ℓ_{2000}	0.46	0.28	0.47	0.31	0.46
c_{300} VS ℓ_{2000}	0.17	0.09	0.18	0.07	0.17
c_{700} VS ℓ_{2000}	0.05	0.05	0.05	0.02	0.05
c_{300} VS ℓ_{1000}	0.11	0.07	0.11	0.02	0.11
c_{300} VS ℓ_{4000}	0.28	0.12	0.30	0.12	0.29
c_{25} VS $\ell_{1000} + g_{1000}$	0.83	0.43	0.86	0.36	0.85
c_{300} VS $\ell_{1000} + g_{1000}$	0.64	0.24	0.67	0.14	0.65
c_{700} VS $\ell_{1000} + g_{1000}$	0.39	0.15	0.41	0.10	0.40
c_{300} VS $\ell_{400} + g_{400}$	0.29	0.07	0.30	0.02	0.30
c_{300} VS $\ell_{2000} + g_{2000}$	0.73	0.33	0.76	0.24	0.74
c_{25} VS g_{1000}	0.89	0.48	0.92	0.43	0.90
c_{100} VS g_{1000}	0.83	0.38	0.86	0.33	0.84
c_{300} VS g_{1000}	0.69	0.24	0.72	0.21	0.70
c_{700} VS g_{1000}	0.41	0.14	0.43	0.10	0.42
c_{300} VS g_{500}	0.38	0.14	0.40	0.10	0.39
c_{300} VS g_{2000}	0.76	0.35	0.80	0.29	0.78
c_{300} VS g_{4000}	0.84	0.43	0.87	0.36	0.85
N	1248	88	1160	42	1206

Notes: Financial aid take-up by aspiration and expectation for college attendance.

Financial Aid Take-Up Behavior In this section, we provide descriptive characteristics for specific types of behavior in the data. That is, we consider the set of individuals who either always accept the cash payment or the financial aid, and document how these students differ from the general population.

Table A7 shows that there are only 37 students who always choose financial aid packages, while 113 individuals always take the cash payment. We find that there are no systematic differences in observed characteristics between the general population and those two specific subgroups.

Table A7: Descriptive Characteristics of Students by Financial Aid Take-Up Behavior

	All	Always		Grant Over Hybrid		
		Cash	F/Aid	25 vs 1,000	300 vs 1,000	700 vs 1,000
Drop-out	0.08	0.09	0.11	0.11	0.12	0.06
High-school	0.24	0.33	0.32	0.19	0.24	0.23
V/College	0.07	0.11	0.05	0.11	0.08	0.07
College	0.61	0.48	0.51	0.58	0.56	0.64
0-20K	0.06	0.05	0.05	0.06	0.08	0.03
20-40K	0.13	0.13	0.11	0.15	0.15	0.15
40-60K	0.24	0.27	0.32	0.22	0.28	0.25
60-80K	0.18	0.23	0.11	0.17	0.13	0.17
80-100K	0.15	0.12	0.22	0.15	0.18	0.14
+100K	0.24	0.19	0.19	0.26	0.18	0.26
Male	0.46	0.62	0.43	0.53	0.52	0.44
Female	0.54	0.38	0.57	0.47	0.47	0.56
Urban	0.80	0.81	0.70	0.85	0.80	0.86
Rural	0.20	0.19	0.30	0.15	0.20	0.14
Quebec	0.30	0.19	0.30	0.18	0.22	0.28
Ontario	0.29	0.18	0.32	0.24	0.34	0.30
Manitoba	0.27	0.34	0.32	0.35	0.28	0.33
Saskatchewan	0.13	0.29	0.05	0.24	0.15	0.09
Numeracy	5.10	4.74	5.24	4.73	4.95	5.13
Full Sample	1248	113	37	89	160	203

Notes: Descriptive characteristics by financial aid take up behavior. Always Cash and F/Aid refers to the subset of students who always chooses the cash and financial aid payment. Grant over Hybrid refers to the students who choose the grant alone but reject the hybrid package.

Then, we analyze the characteristics of students who choose a given grant but reject a similar grant when combined with a loan. While there are small differences between these subgroups and the general population, there is no marked heterogeneity patterns that are consistent across choices. Overall, this table indicates that some choice patterns in the data can not be accounted for using observed heterogeneity. We next analyze whether our measures of risk aversion and time preferences are more predictive of these behaviors.

Table A8: Risk and Time Preferences by Financial Aid Take-Up Behavior

	All	Always		Grant Over Hybrid		
		Cash	F/Aid	25 vs 1,000	300 vs 1,000	700 vs 1,000
(a). Risk Aversion						
Min	-1.60	-0.96	-0.38	-0.03	-0.15	-0.05
1st De.	0.40	0.44	0.40	0.34	0.41	0.39
1st Qu.	0.52	0.60	0.51	0.51	0.52	0.49
40th Cent.	0.60	0.67	0.53	0.61	0.58	0.58
Median	0.64	0.72	0.56	0.66	0.63	0.63
Mean	0.73	0.83	0.74	0.80	0.72	0.69
Sd	1.00	1.26	0.88	0.95	0.66	0.76
60th Cent.	0.68	0.76	0.62	0.72	0.68	0.67
3rd Qu.	0.75	0.82	0.71	0.82	0.78	0.73
9th De.	0.85	0.91	0.82	0.92	0.91	0.85
Max	16.82	10.81	5.48	7.27	5.89	10.67
(b). Time Preferences						
Min	0.00	0.00	0.32	0.00	0.00	0.00
1st De.	0.33	0.00	0.75	0.02	0.42	0.55
1st Qu.	0.72	0.02	0.83	0.47	0.69	0.74
40th Cent.	0.80	0.40	0.91	0.72	0.77	0.80
Median	0.83	0.66	0.92	0.78	0.81	0.84
Mean	0.75	0.50	0.89	0.64	0.73	0.77
Sd	0.27	0.38	0.13	0.32	0.26	0.24
60th Cent.	0.87	0.76	0.95	0.81	0.85	0.87
3rd Qu.	0.91	0.83	0.96	0.89	0.90	0.92
9th De.	0.96	0.90	1.00	0.92	0.95	0.96
Max	1.00	1.00	1.00	1.00	1.00	1.00

Notes: Risk and Time preferences of students by financial aid take up behavior. Always Cash and F/Aid refers to the subset of students who always chooses the cash and financial aid payment. Grant over Hybrid refers to the students who choose the grant alone but reject the hybrid package.

Table A8 shows that individuals who always choose cash over financial aid appear to be more risk averse and less patient than the general population. Conversely, students who always accept financial aid have relatively larger discount factors.

D Parametrization

In this section, we present in detail the specification of background consumption and the future value.

$$
\begin{aligned}
\log(c^k) &= \gamma_{0k} + \gamma_{1k}\mathbb{1}^p_{\text{High-school}} + \gamma_{2k}\mathbb{1}^p_{\text{Voc college}} + \gamma_{3k}\mathbb{1}^p_{\text{College}} \\
&+ \gamma_{4k}\mathbb{1}^p_{\text{20-40K}} + \gamma_{5k}\mathbb{1}^p_{\text{40-60K}} + \gamma_{6k}\mathbb{1}^p_{\text{60-80K}} + \gamma_{7k}\mathbb{1}^p_{\text{80-100K}} + \gamma_{8k}\mathbb{1}^p_{\text{+100K}} \\
&+ \gamma_{9k}\mathbb{1}_{\text{Quebec}} + \gamma_{10k}\mathbb{1}_{\text{Ontario}} + \gamma_{11k}\mathbb{1}_{\text{Saskatchewan}} + \gamma_{12k}\mathbb{1}^p_{\text{Rural}} \\
&+ \gamma_{13k}\mathbb{1}_{\text{Female}} + \gamma_{14k}\mathbb{1}_{\text{Citizenship}} + \gamma_{15k}\mathbb{1}_{\text{Siblings -18 years}} + \gamma_{16k}\mathbb{1}_{\text{Siblings +18 years}}
\end{aligned} \tag{11}
$$

$$
\begin{aligned}
\psi_{iq}(g_q, \ell_q) &= \psi_0 + \psi_1 g + \psi_2 g^2 + \psi_3 \ell + \psi_4 \ell^2 + \psi_5 \ell \times g \\
&+ \psi_6 \beta + \psi_7 \theta + \psi_8 \beta \times g + \psi_9 \theta \times g + \psi_{10} \beta \times \ell + \psi_{11} \theta \times \ell \\
&+ \psi_{12}\text{NU} + \psi_{13}\text{NU}^2 + \psi_{14}\text{NU} \times \beta + \psi_{15}\text{NU} \times \theta + \psi_{16}\text{NU} \times g + \psi_{17}\text{NU} \times \ell \\
&+ \psi_{18}\mathbb{1}^p_{\text{20-40K}} + \psi_{19}\mathbb{1}^p_{\text{40-60K}} + \psi_{20}\mathbb{1}^p_{\text{60-80K}} + \psi_{21}\mathbb{1}^p_{\text{80-100K}} + \psi_{22}\mathbb{1}^p_{\text{+100K}} \\
&+ \psi_{23}\mathbb{1}^p_{\text{High-school}} + \psi_{24}\mathbb{1}^p_{\text{Voc college}} + \psi_{25}\mathbb{1}^p_{\text{College}} + \psi_{26}\mathbb{1}_{\text{Quebec}} + \psi_{27}\mathbb{1}_{\text{Ontario}} + \psi_{28}\mathbb{1}_{\text{Saskatchewan}} \\
&+ \psi_{29}\mathbb{1}_{\text{QC}}\frac{g}{2180} + \psi_{30}\mathbb{1}_{\text{ON}}\frac{g}{5667} + \psi_{31}\mathbb{1}_{\text{SK}}\frac{g}{5064} + \psi_{32}\mathbb{1}_{\text{QC}}\frac{\ell}{2180} + \psi_{33}\mathbb{1}_{\text{ON}}\frac{\ell}{5667} + \psi_{34}\mathbb{1}_{\text{SK}}\frac{\ell}{5064} \\
&+ \psi_{35}\mathbb{1}^p_{\text{20-40K}} \times g + \psi_{36}\mathbb{1}^p_{\text{40-60K}} \times g + \psi_{37}\mathbb{1}^p_{\text{60-80K}} \times g + \psi_{38}\mathbb{1}^p_{\text{80-100K}} \times g + \psi_{39}\mathbb{1}^p_{\text{+100K}} \times g \\
&+ \psi_{40}\mathbb{1}^p_{\text{20-40K}} \times \ell + \psi_{41}\mathbb{1}^p_{\text{40-60K}} \times \ell + \psi_{42}\mathbb{1}^p_{\text{60-80K}} \times \ell + \psi_{43}\mathbb{1}^p_{\text{80-100K}} \times \ell + \psi_{44}\mathbb{1}^p_{\text{+100K}} \times \ell \\
&+ \psi_{45}\mathbb{1}^p_{\text{20-40K}} \times \theta + \psi_{46}\mathbb{1}^p_{\text{40-60K}} \times \theta + \psi_{47}\mathbb{1}^p_{\text{60-80K}} \times \theta + \psi_{48}\mathbb{1}^p_{\text{80-100K}} \times \theta + \psi_{49}\mathbb{1}^p_{\text{+100K}} \times \theta \\
&+ \psi_{50}\mathbb{1}^p_{\text{20-40K}} \times \beta + \psi_{51}\mathbb{1}^p_{\text{40-60K}} \times \beta + \psi_{52}\mathbb{1}^p_{\text{60-80K}} \times \beta + \psi_{53}\mathbb{1}^p_{\text{80-100K}} \times \beta + \psi_{54}\mathbb{1}^p_{\text{+100K}} \times \beta \\
&+ \psi_{55}\mathbb{1}^p_{\text{HS}} \times g + \psi_{56}\mathbb{1}^p_{\text{VC}} \times g + \psi_{57}\mathbb{1}^p_{\text{CO}} \times g + \psi_{58}\mathbb{1}^p_{\text{HS}} \times \ell + \psi_{59}\mathbb{1}^p_{\text{VC}} \times \ell + \psi_{60}\mathbb{1}^p_{\text{CO}} \times \ell \\
&+ \psi_{61}\mathbb{1}^p_{\text{HS}} \times \beta + \psi_{62}\mathbb{1}^p_{\text{VC}} \times \beta + \psi_{63}\mathbb{1}^p_{\text{CO}} \times \beta + \psi_{64}\mathbb{1}^p_{\text{HS}} \times \theta + \psi_{65}\mathbb{1}^p_{\text{VC}} \times \theta + \psi_{66}\mathbb{1}^p_{\text{CO}} \times \theta
\end{aligned} \tag{12}
$$

E Risk aversion, discount factors and background consumption: additional results

In this section, we present additional results on the distribution of risk and time preferences (Table A9), as well as results from linear regressions of the estimated risk and time preference parameters on basic observed characteristics (Table A10). Overall, we find very little evidence that preferences for time and risk are related to parental education and income.

Table A9: Discount factor and risk aversion

	Discount Factor	Risk Aversion
Min	0.00	-1.60
1st Dec.	0.33	0.40
1st Quart.	0.72	0.52
40th Pct.	0.80	0.60
Median	0.83	0.64
Mean	0.75	0.73
SD	0.27	1.00
60th Pct.	0.87	0.68
3rd Quart.	0.91	0.75
9th Dec.	0.96	0.85
Max	1.00	16.82
Correlation	-0.14	

Table A10: Understanding individual preferences

		Regression	
		Discount Factor	Risk Aversion
	Const	0.68***	0.86***
		(0.04)	(0.16)
	Drop-out	Ref.	Ref.
	High-school	0.02	-0.16
		(0.03)	(0.12)
Parental	V/College	0.01	-0.08
Education		(0.04)	(0.15)
	College	0.03	-0.19
		(0.03)	(0.11)
	0-20K	Ref.	Ref.
	20-40K	0.02	0.2
		(0.04)	(0.14)
	40-60K	0.03	0.07
		(0.03)	(0.14)
Parental	60-80K	-0.03	0.1
Income		(0.04)	(0.14)
	80-100K	0.02	0.14
		(0.04)	(0.14)
	+100K	0.03	0.16
		(0.04)	(0.14)
	Rural	0.04*	-0.1
		(0.02)	(0.08)
	Female	0.08***	-0.12*
		(0.01)	(0.06)
	Immigrant	-0.14***	0.19
		(0.03)	(0.12)
	Ontario	Ref.	Ref.
	Quebec	0	-0.05
		(0.02)	(0.08)
	Manitoba	0	-0.06
Province		(0.02)	(0.08)
	Saskatchewan	-0.12***	0.11
		(0.03)	(0.1)
R^2		0.09	0.02
Num. obs		1,248	1,248

Notes: (i) Least squares weighted by the inverse of the standard errors of the estimated individual preference parameters, (ii) standard errors reported in parenthesis, (iii) $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Finally, we report below the estimated coefficients for background consumption associated with the financial aid questions:

Table A11: Determinants of background consumption (financial aid questions)

	Estimates	St. errors
Const	3.71	0.02
High-school	0.53	0.04
V/College	1.01	0.05
College	0.49	0.03
20-40K	-0.71	0.03
40-60K	0.59	0.03
60-80K	0.52	0.03
80-100K	-0.59	0.04
+100K	-0.29	0.04
Urban	0.34	0.03
Male	0.89	0.03
Immigrant	-0.84	0.03
Quebec	0.13	0.03
Ontario	0.96	0.04
Saskatchewan	0.64	0.03
Siblings less than 18	0.07	0.05
Siblings more than 18	0.19	0.02

F Model fit

Table A12: Model Fit

	Predicted Probabilities	Empirical Frequencies
c_{25} VS ℓ_{2000}	0.63	0.46
c_{300} VS ℓ_{2000}	0.19	0.17
c_{700} VS ℓ_{2000}	0.04	0.05
c_{300} VS ℓ_{1000}	0.18	0.11
c_{300} VS ℓ_{4000}	0.28	0.28
c_{25} VS $\ell g_{1000} + g_{1000}$	0.89	0.83
c_{300} VS $\ell g_{1000} + g_{1000}$	0.60	0.64
c_{700} VS $\ell g_{1000} + g_{1000}$	0.24	0.39
c_{300} VS $\ell g_{400} + g_{400}$	0.35	0.29
c_{300} VS $\ell g_{2000} + g_{2000}$	0.74	0.73
c_{25} VS g_{1000}	0.90	0.89
c_{100} VS g_{1000}	0.84	0.83
c_{300} VS g_{1000}	0.64	0.69
c_{700} VS g_{1000}	0.28	0.41
c_{300} VS g_{500}	0.41	0.38
c_{300} VS g_{2000}	0.82	0.76
c_{300} VS g_{4000}	0.78	0.84

Note: c, g and ℓ stand for cash, grant and loan, respectively. For example c_{25} VS g_{1000} refers to the choice between \$25 cash and \$1,000 grant, while c_{25} VS $\ell_{1000} + g_{1000}$ refers to the choice between \$25 cash and the hybrid package (\$1,000 loan and \$1,000 grant).

G Spatial heterogeneity

In Table A13 below, we report the distribution of the willingness-to-pay for a \$1,000 loan, a \$1,000 grant and a hybrid financial aid package containing a \$1,000 loan and a \$1,000 grant, separately for Ontario and Québec. These two provinces, which are the most populated in Canada, are characterized by markedly different levels of tuition fees. Specifically, Québec has low average annual tuition (\$2,180 over the period of interest), while Ontario has significantly higher tuition rates (\$5,667 on average over the period of interest). Québec and Ontario do not only differ in terms of higher education tuition fees, but also in terms of financial aid generosity. The average amount of financial aid offers per student, as well as the fraction of the total financial aid that takes the form

of grants is significantly higher in Québec than in Ontario.For a Québec resident, the median value of securing a $1,000 loan is estimated to be $54. On the other hand, the median young Ontarians, who are faced with higher tuition rates, would pay a higher amount ($71). More generally, the distribution of the willingness-to-pay in Ontario dominates almost everywhere the distribution in Québec. In particular, the first quartile, the median and the third quartile of the distribution of the willingness-to-pay are all significantly higher in Ontario than in Québec at the 5% level. At the third quartile, the values indicating the willingness-to-pay are equal to $137 for Ontario and $100 for Québec.

Students from Ontario also attach higher values to grants as well as hybrid loans than students from Québec. The median Ontarian high school student in our sample is willing to forego $458 to secure a $1,000 grant (45.8 cents per dollar), while the median student from Québec is willing to trade $360 against a $1,000 grant (36 cents per dollar). Differences for hybrid loans are of similar magnitude, with the median willingness-to-pay for a hybrid ($1,000 loan, $1,000 grant) package in Ontario being equal to $432, against $334 in Québec. These findings are consistent with the existence of significantly higher levels of (net) tuition fees in Ontario than in Québec.

Table A13: The distribution of the willingness-to-pay by province

	$c_i^m(\ell_{1000})$		$c_i^m(g_{1000})$		$c_i^m(\ell_{1000}+g_{1000})$	
	QC	ON	QC	ON	QC	ON
1st Dec.	10	8.9	170	228	168	216
1st Quart.	36	31	297	369	274	341
Median	76	68	447	519	409	474
3rd Quart.	133	131	595	708	539	621
9th Dec.	181	183	731	864	636	761

Notes: (i) QC: Quebec, ON: Ontario. ii) Amounts are in Canadian dollars.

H Robustness of the willingness-to-pay

In this section, we provide additional estimation results to assess the robustness of our findings.

H.1 Alternative estimation samples

We first analyze the robustness of the willingness-to-pay for financial aid to the exclusion of impatient students. To do so, we re-estimate the model on the subset of individuals with an (estimated)

discount factor higher than 0.10.[40] This corresponds to 109 agents, which leaves us with 1,139 individuals. Estimation results for this subsample of 1,139 individuals are reported in Table A14. While they indicate a very small upward shift in the distribution of willingness-to-pay for loans, the magnitude of the changes shows that our estimates of the willingness-to-pay are overall robust to the exclusion of impatient students. We do a similar exercise excluding individuals with 2.5% lowest and highest risk aversions. We obtain a median willingness to pay of 62.3$ for a 1,000$ loan to be compared to 60.6$ in our benchmark. Then, we do not use financial aid decisions that involve a cash payment of 700$, for which the assumption that individuals do not smooth consumption is strong. We obtain a median willingness to pay of 65.4$ for a 1,000$ loan to be compared to 60.6$ in our benchmark.

Table A14: Distribution WTP $1,000 loan

	Exclusion of Outliers				Exclusions	
	Discount rate		Risk aversion		700$ cash	
	Loan	Grant	Loan	Grant	Loan	Grant
1st Dec.	6.6	155.9	5.2	123.5	5.0	66.1
1st Quart.	26.3	319.3	24.9	319.2	24.9	289.5
Median	62.6	464.5	62.3	473.2	65.4	436.3
3rd Quart.	118.1	653.0	115.9	660.7	123.1	619.8
9th Dec.	175.6	811.8	170.5	827.0	182.6	776.6

Notes: Distribution of willingness to pay for additional sample definitions. Outliers correspond to the individuals with the highest and lowest 2.5% of risk preferences, and individuals with discount factors lower than 0.10. Choices involving a 700$ cash are not considered in the second specification. Amounts are in Canadian dollars.

Loans questions only We then estimate the model using only financial decisions involving loans. This exercise allows us to alleviate concerns that grant decisions may be driving some of our parameter estimates. We obtain a median willingness to pay of 62.8$ for a 1,000$ loan to be compared to 60.6$ in our benchmark.

[40] While the 0.10 threshold is arbitrary, a descriptive analysis of the reported choices reveals that the large majority of these individuals are "impatient" in the sense that they consistently favor the sooner payment.

Table A15: Distribution of willingness-to-pay (loan questions only)

	$c_m^i(\ell_{1000})$	$c_m^i(\ell_{2000})$
1st Dec.	5.2	4.9
1st Quart.	35.8	34.0
Median	62.8	71.5
3rd Quart.	97.3	124.1
9th Dec.	143.5	187.0

Notes: Distribution of willingness to pay when considering only loan questions. Amounts are in Canadian dollars.

Subset of students expecting and aspiring to attend college Finally, we consider willingness-to-pay for students who expect or aspire to pursue higher education. 93% of our sample expect to pursue higher education, while this share increases to 96% when using data on aspirations. We show that while individuals who expect (or aspire) to pursue higher education tend to have higher valuation of financial aid, the magnitude of the estimates are similar to that of the benchmark model.

Table A16: Distribution of willingness-to-pay (aspiration and expectation)

	Benchmark		Expectation		Aspiration	
	Loan	Grant	Loan	Grant	Loan	Grant
1st Decile	4.6	68.3	5.4	113.4	5.2	85.5
1st Quart.	22.5	293.2	26.3	319.4	25.6	310.7
Median	60.6	449.8	68.9	466.3	67.8	458.8
3rd Quart.	116.7	638.8	128.0	614.3	126.6	601.8
9th Decile	178.6	795.0	174.7	779.9	173.7	769.1

Notes: Amounts are in Canadian dollars.

H.2 Alternative specifications of the future component ψ_{iq}

In this section, we analyze the robustness of our estimates to different specifications of the future value. To that end, we consider a parsimonious specification, where all variables (risk aversion, time preferences, parental income, parental education, region, numeracy score, loan and grant) enter linearly in the future component. Then, we consider an expanded specification, where all variables

are interacted. Finally, we consider another specification where the idiosyncratic shock enters the future component directly.

Table A17: The distribution of willingness-to-pay ($1,000 loans and grants)

	Benchmark		Parsimonious		Expanded		Shock	
	Loan	Grant	Loan	Grant	Loan	Grant	Loan	Grant
1st Dec.	4.6	68.3	5.7	6.5	5.6	67.1	5.7	66.7
1st Quart.	22.5	293.2	24.3	55.1	24.5	314.0	25.8	299.4
Median	60.6	449.8	77.2	158.0	65.0	467.5	66.8	449.6
3rd Quart.	116.7	638.8	156.6	308.1	123.3	662.7	125.0	638.1
9th Dec.	178.6	795.0	293.1	430.4	180.2	831.8	184.3	801.7

Notes: Distribution of willingness to pay for different specifications of the future value. Under the parsimonious specification, all variables (risk aversion, time preferences, parental income, parental education, region, numeracy score, loan, grant) enter linearly in the future components. Under the expanded specification, all variables are interacted. Amounts are in Canadian dollars.

Table A18: The distribution of willingness-to-pay ($2,000 loans and grants)

	Benchmark		Parsimonious		Expanded		Shock	
	Loan	Grant	Loan	Grant	Loan	Grant	Loan	Grant
1st Dec.	5.1	160.9	5.2	27.5	5.8	154.7	5.3	153.4
1st Quart.	29.4	557.7	26.5	166.8	30.3	571.5	32.2	555.0
Median	72.2	792.4	78.5	366.5	75.1	829.4	79.1	801.0
3rd Quart.	129.7	1118.2	158.7	579.5	131.4	1148.9	138.8	1106.9
9th Dec.	182.3	1469.2	286.7	755.2	183.1	1532.5	192.6	1466.5

Notes: Amounts are in Canadian dollars. Under the parsimonious specification, all variables (risk aversion, time preferences, parental income, parental education, region, numeracy score, loan, grant) enter linearly in the future components. Under the expanded specification, all variables are interacted.

H.3 Alternative specifications of background consumption c_0^f

In this section, we analyze the robustness of our estimates to different specifications of the background consumption. We first include the numeracy score into background consumption and then consider an expanded specification where all variables are interacted. Our results are largely unaffected by

changes to the specification of the background consumption.

Table A19: Distribution of willingness-to-pay

	Benchmark		Numeracy included		Interactions	
	Loan	Grant	Loan	Grant	Loan	Grant
1st Decile	5.1	66.5	4.4	66.7	5.1	65.9
1st Quart.	23.6	282.6	24.2	287.6	23.5	283.1
Median	66.8	449.3	60.4	433.4	66.3	450.8
3rd Quart.	132.7	638.9	116.5	617.4	130.9	638.5
9th Decile	204.2	832.3	175.3	770.3	203.4	830.2

Notes: Amounts are in Canadian dollars. All interactions refer to a specification where all variables (parental education and income, gender, citizenship, province, family size) are interacted with each other.

I Determinants of higher education expectation and aspiration

In this section, we attempt to replicate the findings of Belley et al. (2014). Since we do not have access to actual college attendance data, we use as a proxy data our measures of higher education expectation and aspiration. Table A20 reports the determinants of higher education attendance expectation and aspiration. Specifications (1) and (2) are very close to that of Belley et al. (2014) (Table 3, page 674), although we have fewer controls for family background. Specifications (3) and (4) include our measures for time and risk preferences. Our results are in line with Belley et al. (2014) regarding the role of gender and education. However, we do not find any effect of the immigration status of the student, which is probably explained to some degree at least by the very small number of immigrants in our sample. We also show that parental education matters in the sense that students with a least one college educated parent are more likely to aspire and expect to enroll in higher education.

Table A20: Probit models of higher education expectation and aspiration

		Regressions			
		Expectations		Aspirations	
		(1)	(2)	(3)	(4)
	Const	1.076**	0.97**	1.755***	1.611**
		(0.357)	(0.375)	(0.506)	(0.52)
	Female	0.413***	0.343**	0.503**	0.45**
		(0.116)	(0.119)	(0.154)	(0.159)
	Immigrant	0.096	0.004	-0.442	-0.492
		(0.225)	(0.23)	(0.395)	(0.393)
	Rural	0.268	0.182	0.164	0.064
		(0.157)	(0.158)	(0.194)	(0.197)
Parental	High-school	0.263	0.21	0.232	0.183
Education		(0.186)	(0.19)	(0.232)	(0.239)
	V/College	0.091	0.067	-0.077	-0.136
		(0.237)	(0.243)	(0.286)	(0.292)
	College	0.649***	0.591**	0.55*	0.497*
		(0.182)	(0.187)	(0.228)	(0.235)
	20-40K	-0.605	-0.598	-0.351	-0.347
		(0.312)	(0.318)	(0.367)	(0.378)
Parental	40-60K	-0.503	-0.52	-0.22	-0.248
Income		(0.303)	(0.308)	(0.357)	(0.367)
	60-80K	-0.41	-0.395	-0.061	-0.051
		(0.319)	(0.323)	(0.38)	(0.39)
	80-100K	-0.364	-0.372	-0.039	-0.064
		(0.323)	(0.328)	(0.386)	(0.397)
	+100K	0.02	0.027	0.224	0.211
		(0.332)	(0.34)	(0.397)	(0.41)
	Siblings less than -18	0.179	0.173	0.617	0.61
		(0.206)	(0.218)	(0.395)	(0.421)
	Numeracy	0.041	-0.021	-0.093	-0.171
		(0.162)	(0.164)	(0.2)	(0.203)
	β		0.559**		0.566*
			(0.19)		(0.236)
	θ		-0.112**		-0.094*
			(0.04)		(0.045)

Notes: Notes: (i) standard errors reported in parentheses, iii) $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.