Alma Mater Studiorum – Università di Bologna

DOTTORATO DI RICERCA IN

ECONOMIA

Ciclo XXIX

Settore Concorsuale di afferenza: 13/A1

Settore Scientifico disciplinare: SECS-P/01

The Determinants of Individual Performance: Empirical Essays on the Importance of Soft Skills and Monetary Incentives

Presentata da: **VERONICA RATTINI**

Coordinatore Dottorato **Relatore**

VINCENZO DENICOLÒ **MARIA BIGONI**

Esame finale anno 2017

Abstract

This dissertation consists of three papers. The first paper "Managing the Workload: an Experiment on Individual Decision Making and Performance" experimentally investigates how decision-making in workload management affects individual performance. I designed a laboratory experiment in order to exogenously manipulate the schedule of work faced by each subject and to identify its impact on final performance. Through the *mouse click-tracking* technique, I also collected interesting *behavioral* measures on organizational skills. I found that a non-negligible share of individuals performs better under externally imposed schedules than in the unconstrained case. However, such constraints are detrimental for those good in self-organizing. The second chapter, "On the allocation of effort with multiple tasks and piecewise monotonic hazard function", tests the optimality of a scheduling model, proposed in a different literature, for the decisional problem faced in the experiment. Under specific assumptions, I find that such model identifies what would be the optimal scheduling of the tasks in the Admission Test. The third paper "The Effects of Scholarships and Tuition Fees Discounts on Students' Performances: Which Monetary Incentives work Better?" explores how different levels of monetary incentives affect the achievement of students in tertiary education. I used a *Regression Discontinuity Design* to exploit the assignment of different monetary incentives, to study the effects of such liquidity provision on performance outcomes, *ceteris paribus*. The results show that a monetary increase in the scholarships generates no effect on performance since the achievements of the recipients are all centered near the requirements for non-returning the benefit. Secondly, students, who are actually paying some share of the total cost of college attendance, surprisingly, perform better than those whose cost is completely subsidized. A lower benefit, relatively to a higher aid, it motivates students to finish early and not to suffer the extra-cost of a delayed graduation.

Contents

List of Figures

List of Tables

Chapter 1

Managing the Workload: an Experiment on Individual Decision Making and Performance

1.1 Introduction

This paper examines how individuals manage their workload when they have several tasks to solve in a limited timespan. In particular, the aim of the study is to understand if people are effective when choosing how to allocate their effort across tasks through time, or whether better performance could be achieved by restricting the individual's decision space.

Typically, economic models of individual production have assumed that agents choose optimally the cognitive operations for each working domain. However, the evidence described in some top-selling managerial books [Allen (2015), Covey (1991) and Crenshaw (2008)] and in the recent cognitive literature, Crawford and Wiers (2001), recognizes that decisions on workload management require cognitive operations, such as information search and processing, which involve costly and complex valuations, and agents often fail to choose such operations optimally.

The present project is interested in studying the individual decision-making process insofar as it affects productivity in a context of workload management. The fact that individuals may adopt simple heuristics in work organization, is an interesting feature of human behavior *per se*. Nonetheless, understanding the relation between the organizational strategies and the individual performance is essential for the economics of production. The current study experimentally evaluates subjects' behavior in the framework offered by the "Admission Test" to the faculty of Economics of the University of Bologna. In this context, individuals have to solve several tasks in a given timespan. In the baseline condition, students answer the test without any restrictions on their decision space, therefore deciding freely how to allocate the time across tasks and the sequence in which they want to answer. Two treatments are designed to test if restrictions on such space affect their final performance. In such cases, the test will be administered to subjects either with a fixed sequence of questions, or with a fixed time for completing each task.

The application to a field context adds interest in terms of policy evaluation for the case of study. In particular, the replication of the "Admission Test" of the University of Bologna will give the opportunity to test whether the design of such an assessment method has an impact on the selection of the students. In particular, since the current version of the test restricts both the timespan and the sequence of the sections, which are thematically clustered questions, this study will test whether such restrictions affect the students' performances and which type of individuals benefit or are at a disadvantage from such a test format.

Finally, the contribution of the present paper relies also on the adoption of a rich data-gathering method. A sophisticated software design has allowed the collection of comprehensive information on individual behavior using the *mouse-tracking* technique and, therefore, a complete map of the solving strategy adopted by each student is revealed. Such information is indeed used to provide new insight into subjects' behavior and for refining the empirical analysis.

Results show that, on average, subjects performed better when they could freely choose how to manage their task-load. However, treatment restrictions significantly improve the performance of those individuals who have shown inefficient organizational skills. Such heterogeneity demonstrates that the sample composition would be crucial for interpreting the aggregate results in similar contexts of study.

The paper proceeds as follows: Section 1.2 revises the literature of interest, while section 1.3 describes the experimental framework and summarizes the theoretical predictions for such decisional context. Section 1.4 presents the design and the procedure of the experiment. Section 1.5 summarizes the results of the experiment and describes the robustness checks performed. Finally, section 1.6 offers concluding observations.

1.2 Literature Review

The importance of individual decision making in Economics has been emphasized since the breakthrough work of the Nobel laureate, Herbert Simon, who showed how the decision-making process of individuals might be rationally bounded in several domains [Simon (1955), Simon and Chase (1988)] and clearly pointed out the necessity of characterizing such a process, rather than relying on the Homo Economicus assumptions.

Even though, in recent decades economists started to take into consideration the findings from the psychological field related to the bounded rationality of individuals [for a review DellaVigna (2009)], there is still a lack of extensive studies in workload management and effort allocation. Traditionally, labour economists have defined effort as "the pace or intensity of work" [Johnson (1990)], and they have usually ascribed such input to individual rational choice, practically governed by economic incentives. However, recent studies on effort allocation and task prioritization have shown that subjects may not be so effective in choosing how to manage their workloads, even when they are incentivized to do so. We will now present and discuss some of those studies that have focused on such decisional domains directly related to the present paper.

It has been shown both empirically and theoretically that, when agents are left free to manage their work independently, real costs in term of performance emerge with respect to the potential output obtained under an imposed schedule of work. The paper most closely related to the present research is Buser and Peter (2012). The authors have examined, in a laboratory experiment, the effects of multitasking on the performance of subjects in two games: a word-search puzzle and Sudoku. Their results show that, when subjects are left free to to organize themselves independently in the two games, they perform worse than those who face a given sequence of the two. Moreover, the evidence has suggested that forcing individuals to multitask is detrimental with respect to the sequential rationale. However, in this experiment, decisions on the time allocation are restricted and the duration of tasks are always imposed by the researchers. Given that this choice dimension is often encountered in problems of work division, and it could significantly influence the overall individual performance, this paper extends the analysis of Buser and Peter (2012), by considering contexts where time-management is not restricted. In addition, the conclusions provided by the same authors would be valid for the two-tasks case; the present research opens the set-up to the multiple-tasks framework, further enriching the analysis of such topics.

The evidence described in Buser and Peter (2012), that working in parallel on more than one task is detrimental with respect to a sequential working rationale, is further confirmed in the field by the empirical study of Coviello et al. (2015) and by the related theoretical framework in Coviello et al. (2014). In particular, the authors have analysed the work schedule of the Italian judges of the Labour Court of Milan. Judges who work in this Court, exogenously receive a stream of cases that they have to undertake and, as underlined by the authors, they usually work in parallel on more than one assigned project. The evidence has proved that such organization of case-load has a substantial negative impact on the productivity of the Italian judicial system. Moreover, in Coviello et al. (2014) the authors propose a model of dynamic production to prove that the optimal strategy of workload management in such domain would indeed follow a sequential schedule. In their model, the effort allocation is governed by the input rate, i.e the rate at which each task is began, and it is taken as exogenous or determined by co-workers. The present paper complements and extend their analysis by studying workload management in a context where such rate would be endogenously chosen by the individual.

As previously anticipated, limits on agent decisional process not only emerge for choices on the sequence of work, but also for choices regarding the time allocation among several tasks. Tice and Baumeister (1997) show in two longitudinal studies that students identified as "procrastinators", obtained lower grades than those students who do not procrastinate. In principle, there should not be any difference in performance in relation to whether the task is done far ahead of the deadline or just in time. However, the authors find significant differences in performance between the two types of students. As the authors suggest, procrastinators may outweigh the short-term benefit of the postponement and underweight the long-term cost of delaying, allocating, therefore, sub-optimal effort over time; or it may be the case that they simply and naively believe that such a postponement would improve their performance, while it actually constrains their available time. In addition, it could be also that those who procrastinate lack a clear strategy for prioritizing tasks over time; that is, they have poor time-management skills. Which of these proposed explanations is more likely is still unclear and the present research aims to provide new insight into such question.

The above result is further confirmed by Ariely and Wertenbroch (2002), who proved in a field experiment that when students have to work on several tasks without any strict deadlines per task, they perform significantly worse than those who self-organize their effort through costly deadlines. Moreover, internally imposed deadline are not as effective as external ones in improving task performance. Such experimental findings confirm more robustly the evidence that individuals may not choose effectively the time allocation that would maximize their performance. In addition, the study shows that limiting subjects' decision space by externally imposing a fixed time per task, could be an effective way to raise overall output. Contrasting evidence emerges from a filed study by De Paola and Gioia (2016), the authors find that by imposing a more binding deadline on students sitting intermediate exams generates negative effect on grades with respect to the group of students who face less time pressure. Further analyses indicate that such results are actually driven by female students.

Notice, however, that, with the exception of Buser and Peter (2012) and De Paola and Gioia (2016), the evidence discussed above focuses on tasks to be completed in a long time period. In order to limit future uncertainty regarding available time for task completion, the present study focuses on shorter tasks and on a context of complete information on the total available time. Moreover, important innovations of the present research rely also on the adoption of an advanced data gathering method, the *Mouse-tracking* technique.[1] This method consists of tracking the mouse-click of each subject during the experiment and recording the sequence and the duration of each step in the proceeding working strategy. Due to its completeness and its richness in information acquisition, this method has been widely adopted both by the psychological and economic literature [see for example Arieli et al. (2011), Gabaix et al. (2006) or Johnson et al. (2002)]. By applying such a sophisticated methodology, this study could extend the previous results of the literature, by describing the choices, the sequence and the duration

[1] Firstly introduced in Johnson et al. (1989).

related to each mouse-click and the implicated use of time, and by estimating the treatment effects for different types of organization chosen by subjects in the experiment.

To conclude, the purpose of the present study is to improve our knowledge on the understanding of individual decision making in workload management, by providing evidence from a laboratory experiment, which relies on a field framework where such a choice domain is present. In particular, the current research aims to investigate further the importance of both the time allocation and of the choice of the sequence of work in a context with multiple tasks and with complete information on the total available time, refining and extending the findings raised by the above reviewed literature.

1.3 The Framework

1.3.1 The TOLC Test

The experiment is based on the format and on the tasks proposed by the electronic version of the Admission Test to the Faculty of Economics at the University of Bologna.

The test is called "TOLC" and is provided by a private agency, i.e CISIA [2]. This test was adopted by the School of Economics, Management, and Statistics by the University of Bologna in 2013, changing from a paper-based format to the CISIA electronic version. The importance of such implementation is not only related to its implications for the admission of candidates, but also for its extensive coverage among Italian universities. In 2016, in fact, about 56 universities have decided to adopt such test, namely almost 66% of the national total [3].

The TOLC test is composed by three sections: logic, verbal comprehension and mathematics. Sections are clusters of questions that are focused on the same field of knowledge. The total number of questions is 36, divided into 13 questions of logic, 13 of mathematics and 10 tasks of verbal comprehension. The tasks have a multiple-choices format and there is only one correct answer. The total available time to answer all the questions is one hour and thirty minutes, which is uniformly distributed across sections, i.e 30 minutes per section (not combinable). Moreover, students must follow the order of the sections imposed by the format: logic, verbal comprehension and mathematics. This implies that switching from one section to the others is irreversible, so it is not possible to return to the previous thematic area.

As explicitly stated in the instructions provided by the agency, the logic and the verbal-comprehension tasks do not require previous training or particular skills. The mathematical questions cover the program encountered during the first four years of high school. The questions in the verbal-comprehension section are related to two different essays which students have to read in order to answer, always having the possibility to return to the text. Questions

[2]For more information: http://www.cisiaonline.it/area-tematica-tolc-cisia/home-tolc-generale/

[3]List from the University of Bologna.

in the logic section are intended to test the concentration aptitudes of the candidates and their ability in inferring the conclusions from precise statements.

In the figure below, a screen-shot of the test is shown in order to better clarify the structure of the "Admission Test". From figure 1.1, it is possible to distinguish: A) the three sections' buttons; B) the text of the first question with the related answer options; C) the buttons corresponding to the other questions in the section; D) the time bar, which scrolls down as the time goes by.

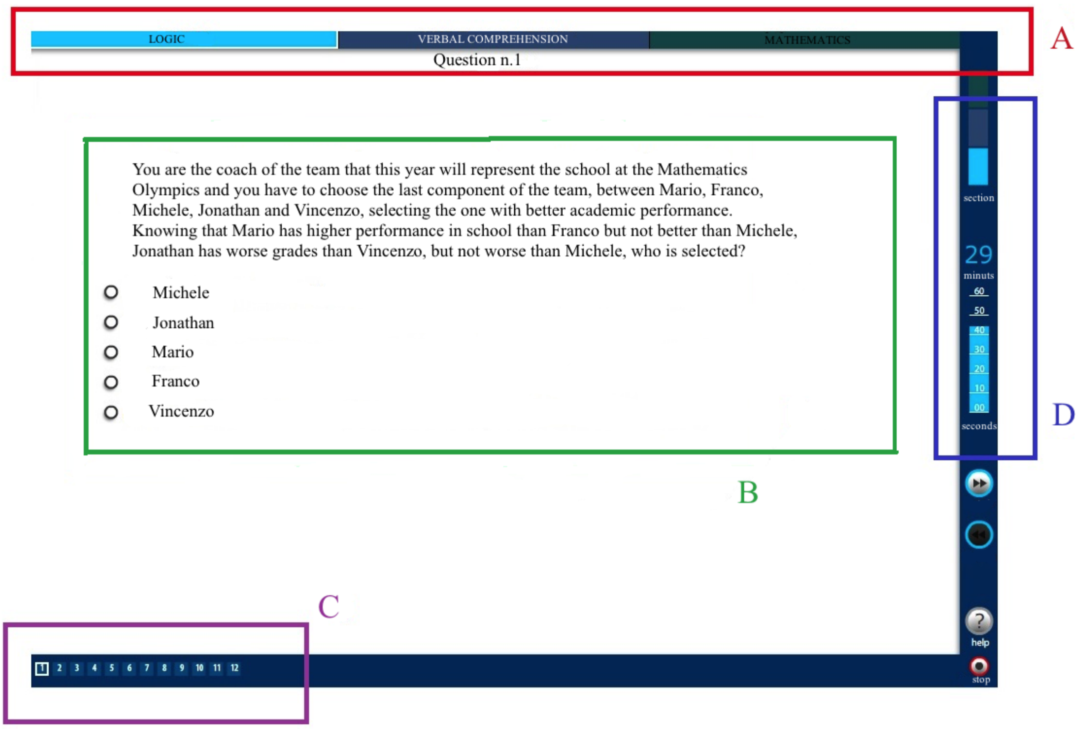

Figure 1.1: The Admission Test

The above description of the test's format aims to underline that the "Admission test" offers a good opportunity to answer the research questions of interest. In particular in this context students face a problem of workload management, since they have to solve several questions in limited available time, with the possibility to choose how to allocate time across tasks and in which sequence they want to answer. Moreover, differently form previous studies, here there is no uncertainty about the total available time, questions are short and there are more than two tasks. For all the above reasons, the TOLC test would be an ideal set-up for studying decisional problem in workload management and, indeed, to extend and refine previous findings of the literature.

This field framework is also interesting due to the possibility of retrieving evidence useful for the university. Understanding how such new implementation may have impacted the performance of students could have relevant policy implications for the design of the student-selection mechanism.

Finally, the electronic format of the TOLC has also brought the advantage of replicating such environment into the laboratory without inducing any framing effects and overcoming the typical difficulties of recreating the field conditions in laboratory experiments.

1.3.2 The Experimental Framework

The experimental framework is based on the TOLC test format previously described. In particular, students have to answer a similar number of questions, 34 tasks, which are distributed across the three sections as in the field case. Moreover, no changes are adopted in respect of the total available time, which would be exactly 1 hour and 30 minutes.

The important feature of the experimental version of the test relies on the data collection. In particular, since the objective of the study is to understand how individuals manage their workload, it is fundamental that the researcher observes all the steps that have brought the subjects to take each action. In order to meet the above objectives, I adopted the *mouse-tracking* technique. The mouse-tracking method was firstly used by Payne et al. (1993) and now it has been accredited by a wide set of scientific domains. As previously described, its implementation on this framework has not only the advantage of mapping completely the decision process of individuals, but also of retrieving interesting and unique measures, principally helpful in setting the difference between the present research and previous literature.

The replication of the test and the implementation of the tracking technique were performed using the Ztree software, Fischbacher (2007). In particular, during the experiment, all subjects enter a computer laboratory where they have an isolated workstation with a laptop and a connected mouse. When the test starts, the Ztree software begins to record all the information relating to the click of the mouse used by each experimental subject. In particular, measures on the timing and on the space where the click is located are recorded in the database. This information allows the recording of the time spent by each subject on answering specific questions, the number of different answers to the same question, the adopted sequence of work, or the times they looked at the same question before answering, etc.

The picture below shows a screen-shot of the terminal during the experimental session. It should be noted that the Ztree software allows for perfect replication of the TOLC test framing, since both the structure and the interface were almost equal to the the field case: see figure 1.1.

The main distinction between the experimental framework and the field case relies on the type of incentives that students face. In particular, while the TOLC test has implications for students' admission to the undergraduate program in Economics at the University of Bologna, the experimental replication provides monetary incentives. Notice, however, that, even if the rewards are different, the payment scheme adopted in the experiment perfectly replicates the TOLC evaluation's scheme. Students, in fact, receive 8 euro as a show-up fee, and the residual

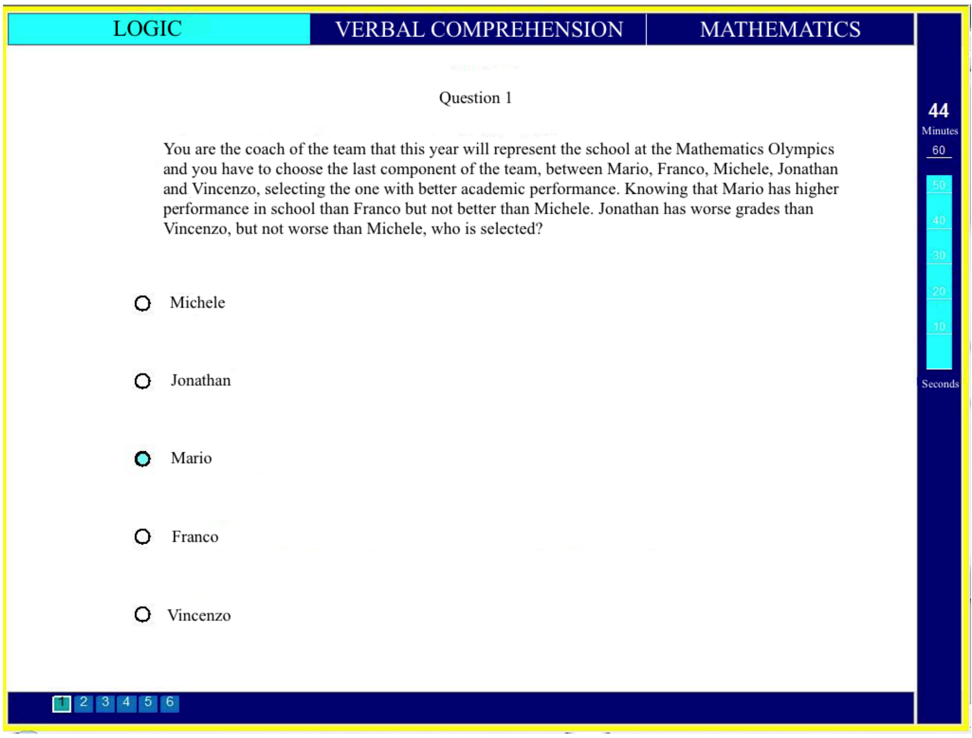

Figure 1.2: Ztree Replication

part of their earnings are calculated as in the field case: 1 point for each right answer, 0 points for missing and -0,25 for incorrect answers. The final score will be expressed in points, which are the Experimental Currency Units and they will be converted into euro at the following rate: 2 ECU = 1 euro. This conversion rate implies that the maximum payment would be 25 euro while the minimum is 3.75 [4].

1.3.3 The Theoretical Predictions

In this framework, students have to decide how to allocate their effort among several tasks, in order to maximize their final payoff, which is given by the sum of the rewards obtained by completing each task. Specifically, at each point in time, each individual has to decide which task should be processed first (*sequence of work*) and how long they should stay on each question (*time allocation*). Typically, this problem is known in the literature as the "*Exploration-Exploitation Dilemma*", where a subjects has to decide how much she wants to search for the best option (*Exploration*) and how much she wants to exploit the current choice (*Exploitation*), given that she has several alternatives to select from (all giving the same reward on completion) and that each option has a random completion time. The so called Gittins Index is the dynamic allocation index that was firstly proposed in Gittins and Jones (1979) to solve the above dilemma.

[4]These amounts include also the 8€show-up fee. The fee has been chosen in order to more than cover the eventual loss from answering all the 34 questions incorrectly, the minimum earnings is, in fact, 3.75€.

It was initially presented for the "Multi-Armed Bandit" process and then for the "Job Scheduling" problem. In the following paragraphs, the attention is focused on the scheduling problem, given that is the dilemma that most closely resembles the present decisional domain.

In the "Job Scheduling" problem, each agent has a queue of jobs to process and she has to decide how to organize the work in order to minimize the mean number of remaining jobs in the queue. Each job yields a positive reward on completion, which, in turns, it happens at random times. Specifically, define as $Ve^{-\delta t}$ the present value of the reward obtained on completion at time t of a generic job. Denote by $F(\cdot)$ the distribution of the service time requested by each job to be completed and by $f(\cdot)$ its density.

The Gittins Index for each job is given by:

$$\nu(x) = \sup_{t>x} \frac{V\int_x^t f(s)e^{-\delta s}ds}{\int_x^t [1-F(s)]e^{-\delta s}ds} \tag{1.1}$$

where the index of the job, in state x, represents the probability that the job will be completed within $t-x$ of additional service time, given that it was not completed before, times the reward on completion V.

As explained in Gittins and Jones (1979), the optimal policy is to process, at each point in time, the jobs in order of decreasing indexes $\nu(x)$. In addition, the authors showed that such a policy would be highly simplified for specific families of service time distributions. In particular, they looked at the hazard rate of each job, which is defined as $\rho(s) = f(s)/[1-F(s)]$. If the jobs belong to the *Decreasing Hazard Rate - DHR* class, where the probability of completion decreases as s increases, it would be optimal to process them in order of least attained service, i.e *Foreground-Background (FB)* policy. In contrast, if the service time distribution belongs to New Better than Used in Expectation - NBUE, as for the case of *Increasing Hazard Rate - IHR*, the *First-Come-First-Served (FCFS)* discipline minimizes the number of jobs in the queue.

Aalto et al. (2009) extends this framework to the class of non-monotonic hazard rate and discusses how the optimal policy would change in such cases. In particular, they consider the case where the hazard rate is piecewise monotonic (firstly increasing and then decreasing), and they showed that the optimal policy would be FCFS+FB(θ^*). Specifically, the optimal schedule would be to serve, without switching, the jobs until their attained service reaches θ^*. Jobs with attained service greater than θ^* are served according to FB, and jobs with no attained service have priority over those with attained service of at least θ^*. The proof relies on the assumption of no-crossing between job hazard rates and on the fact that each job receives attention as long as its hazard is increasing, i.e each job receives its own optimal quantum of service θ^*. After this optimal quantum has been reached and each job hazard has started to decrease, jobs are processed in order of least attained service (FB), starting therefore from the one which has received less service time to the one who has been worked the most.

Such predictions imply that the optimal policy would depend on the distribution of the job hazard rate. In this work, I leave this unrestricted and I will explore empirically what the data tells about the hazard distribution of this kind of jobs, and I will discuss the related implications in terms of optimal behavior in section 1.5.

1.4 Experimental Design and Data

1.4.1 Treatments

The present experimental study is divided into three stages. During the first baseline stage, all the subjects face half of the total questions, i.e 17 tasks, whereof six are from logic, five from verbal comprehension and six from mathematics, with half of the total available time, i.e 45 minutes. Notice that this distribution of questions is in complete accordance with the field case, with the exception that I have excluded two questions to maintain symmetry across stages. In this stage, students proceed to solve the test without having any constraints either on the sequence or on the time per question. For the sake of simplicity, I will define this first stage as “Unconstrained”.

After this first part has been completed, subjects will be randomly allocated to three treatments: “Unconstrained”, “Fixed Time” and “Fixed Sequence”. In the first treatment, subjects answer the remaining 17 questions of the test, again without any restrictions, as in the first stage. In the Fixed Time treatment, they instead answer the remaining questions with a given time per question, which is calculated by dividing the total available time by the number of questions, i.e 40 minutes / 17 tasks = 2 minutes and 22 seconds. Notice that, in this calculation, five minutes were dropped, since this was the maximum time available for reading the text in the verbal comprehension part (this upper bound was computed by considering the registered maximum value from the pilot session). Moreover, in this treatment, when subjects change question, the timer for the switched question would stop and it would re-start as soon as the subject returned to that task.

Finally, a third group of subjects is allocated to the Fixed Sequence treatment. In this case, subjects answer the same remaining 17 questions of the test following a given sequence of the sections, the same as the one provided during the actual admission, i.e first logic, then verbal comprehension and finally mathematics. Moreover, within each section, subjects have to follow a given sequence of questions, following the ascending numeration of the tasks (question 1, then 2, then 3, etc...). Notice that, since the Fixed Sequence treatment has the purpose of denying the subjects the opportunity to choose the answering sequence, if subjects switch questions, they will not have the possibility to come back to the switched task again (the button related to the changed question would disappear from the screen).

Finally, in the third stage of the experiment, subjects answer a general questionnaire that collects information on gender, age, residence, risk-preference and impulsivity [measured by the “Cognitive Reflection Test”, see Frederick (2005)]. After having completed this final stage, subjects are paid and then leave the experimental session. In Appendix B, the instructions, a sample of questions from each section and the final Questionnaire are reported.

The design is summarized in the following table.

Table 1.1: Design

	Baseline Treatment	Sequence Treatment	Time Treatment
Stage 1	Unconstrained	Unconstrained	Unconstrained
Stage 2	Unconstrained	Fixed Sequence	Fixed Time
Stage 3	Questionnaire	Questionnaire	Questionnaire

Notice that, contrary to the TOLC, the baseline format of the test implies that students do not have a fixed time or a fixed sequence for each section. The reason for this discrepancy relates to the aim of the current research. In particular, the main focus of this study is to understand how individuals *freely* behave in such context, without constraining any aspect of their choice space. For this reason, not imposing any time or order restrictions during the baseline treatment of the experiment, allows us to observe unconditional behavior and specifically to test what are the effects of imposing such bounds on performance. In addition, the findings from this baseline design help in understanding the effects of the current format restrictions of the TOLC test, with respect to the previous unconstrained paper-based test, on applicants' performance and on their admission into universities.

1.4.2 Procedures

In the present section I will describe the composition of the subject pool and the details related to the experimental sessions.

For the participation in this study, I recruited students currently enrolled at the fifth-year of high-school. At this grade, students might decide to apply for university in few months representing, therefore, a potentially highly interested subjects pool. The recruitment of such pool was conducted in an Italian high-school, located in Bologna, named "Liceo Classico Statale Marco Minghetti".

In accordance with the director of the school, all the students enrolled at the fifth grade by October 2015 were allowed to participate in the present study. The experiment took place at the BLESS Laboratory (Bologna Laboratory for Experiments in Social Science) of the Department of Economics in Bologna from 24th of November to 5th of December 2015.

The recruitment was done in two phases. Firstly, for each enrolled student I collected a form, where contacts (personal e-mail, mobile and home telephone number, address of residence) were provided. In this occasion, students have also to state the preferred mean of communication for receiving details about the date and place of their future participation into the study. After having recorded the above information, students where randomly allocated to treatments. Finally, for each treatment, three dates where offered for participation (each option befall in a different day of the week: Monday, Tuesday, etc.). With such procedure, each student received invitation according to the preferred mean of communication and he/she was asked to choose one of the proposed dates [5].

The experiment started with an introduction explaining the rules of the first stage and with a practising example. Subjects stay in the laboratory for approximately 1 hour and 45 minutes and the average payment was of about 15,30 euro, which is in line with the average salary potentially earned for the same timespan by an high-school student in Italy.

Finally, it is important to notice that the choice of this specific subject pool has brought some important advantages. In particular, the fact that such students will apply to University in few months has probably pushed their participation by means of their intrinsic motivation for experiencing the test. In this way, the standard problem of the saliency of incentives in laboratory is potentially reduced. Moreover, active participation in this study was also enhanced by the fact that typically the training for such kind of test is provided by manuals and books which offer just indicative exercises. While, in this replication, subjects could face the actual difficulty of the TOLC test, since the proposed questions are exactly the ones of the last admission wave.

For the above reasons, even if subjects are incentivized through monetary rewards (following a pay-for-performance scheme), the time at which they are recruited and the type of questions used, suggest that students' effort is further enhanced by means of their intrinsic motives.

1.5 Results

Section 1.3.3 presented a general model used in the literature to address the Job Scheduling from an optimal perspective.

In this section I will rely on the predictions of the model to discuss what are the expected treatments' effects. The details of the model and formal proofs are presented in Chapter 2.

As in the Job Scheduling problem, in this experiment subjects are incentivized to maximize the sum of their final rewards and, at each point in time, they have to decide which task they want to undertake and for how long they want to work on that task to maximize their payoff. As described in Chapter 2, the optimal policy would depend on the characteristics of the distribution of the hazard rate of the tasks under consideration and thanks to the data gathered through the *mouse-tracking* technique, I show that, for the experimental tasks, the hazards are piecewise monotonic: firstly increasing and the decreasing. Notice that, if one considers the type of questions included in the Admission Test, it does indeed make sense to think that this is the case. At first, subjects have to carefully read the question and think about the answer. If

[5] The participation rate was of about 65% since 87 subjects participate, out of 134, in one of the 9 sessions.

they spend to little time on a question, they run the risk of misinterpreting it and make mistakes even if they knew the answer. But then, if they dont know the answer at first, it is unlikely that they will find it out by spending more time on the question.

Such characterization of the *hazard rate* in our context will help in understanding what the optimal policy would be for the "Task Scheduling" problem. In particular, the results from Aalto et al. (2009) suggest that students should follow a FCFS+FB(θ^*) policy, where θ^* represents the critical amount of processing time, i.e the optimal quantum of time that the agent should allocate to each task. Notice that such an optimum would eventually change according to the specific hazard rate of each task and for each individual. In this sense, I do not exclude the possibility that the hazard rate might be heterogeneous, both within subjects and across tasks and across subjects within question. However, since I assume no-crossing between individual specific hazards, the Gittins rule, which assigns specific (optimal) quantum of time to each task, will still define what the optimal (individual) strategy should be.

Chapter 2 firstly confirm the optimality of the FCFS+FB(θ^*) policy in the present experimental context and it describes in more details how the such policy would be applied in such decisional framework.

The prediction that follows from the optimality of FCFS+FB(θ^*) policy, is that the average score obtained by those subjects who have faced the second part of the test under the two treatments' conditions is expected to be lower than the score obtained by those who have faced the test without any restrictions. Such a result arises from the fact that, when students can freely organize their work, they should apply the optimal FCFS+FB(θ^*) policy, meaning that they should select and prioritize tasks over time, postponing the difficult questions to the end of the test and allocating to each task the optimal quantum of service time, which by no means has to be equal to the average time per task. For these reasons, I expect that both the constraints imposed by the treatments will prevent subjects from following this strategy and, therefore, will reduce their performance levels.

In the following section, I proceed by firstly presenting summary information on the subject pool and then discussing the treatments' effects.

1.5.1 Difference-in-Differences

The sample consists of 87 subjects. As explained in section 1.4, subjects were randomly allocated to three treatments. The summary statistics and the sample balancing is shown in tables A.1 and A.2 of the Appendix A. Table A.2 shows that there are about 30 subjects per treatment group and, thanks to the random allocation, there are no significant differences in all the relevant measured characteristics among the three treatments' groups. In particular, even if different percentages of individuals have taken part in some economic courses across the three groups, the Wilcoxon-Mann-Whitney test fails to reject the null hypothesis that the three samples come from the same population[6]. The same hypothesis is not rejected when I test for difference in risk attitude, in impulsive behavior and in the ability of subjects across groups.

This confirms that the allocation of subjects into the three groups was random in respect of all the above characteristics.

In the following analysis, I will present results of treatments' effects on performance, by looking at the score obtained in solving the tasks of the "Admission Test". In particular, the output of interest is computed by summing up all the points for correct answers and by subtracting the total penalty for the wrong ones.

Table 1.2: Average Score for Baseline Treatment and Fixed Sequence Treatment groups

	Part 1			Part 2			
	Baseline	Fixed Sequence		Baseline	Fixed Sequence		
Outcome	Unconstrained	Unconstrained	Δ_1	Unconstrained	Fixed Sequence	Δ_2	DID
Score	10.091	9.741	-0.351	10.667	10.205	-0.462	-0.111
Std. Error	0.684	0.444	0.778	0.684	0.551	0.778	1.100
t			-0.45			-0.59	-0.10
P>t			0.653			0.554	0.920

Notes: *Score* is the outcome variable, which sums the points obtained from the correct answers minus the penalty encountered from the wrong answers. Δ_t represents the average difference in the scores, in each part t, faced by the subjects in the groups "Baseline" and "Fixed Sequence". *DID*: represents the Difference-in-Differences estimate, which is given by $\Delta_t - \Delta_{t-1}$. Legend: On the left, results related to the fist part of the test, $Part_1$. On the right, results related to the second part of the test, $Part_2$. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$

In table 1.2, the average score obtained in the two parts of the test is shown both for the Baseline Treatment group and for the Fixed Sequence Treatment group.

From table 1.2, it is possible to see that, statistically, there are no significant differences among the two groups in the scores for part 1 of the test, suggesting that the two groups do not differ in terms of "ability". In the second part of the test, students from the Fixed Sequence Treatment have to solve the test through facing a fixed sequence of the questions, as explained in detail in section 1.4. Even in the second part of the test, the two groups obtained similar results suggesting that, in this context, solving the test through a sequential rationale does not affect individual performance.

In table 1.3, the results of the Fixed Time Treatment group are presented.

As previously explained, the Fixed Time Treatment group has to solve the first part of the test with no restrictions, while, in the second part, students in this group are forced to answer all the questions with a given time per task, i.e 2 minutes and 22 seconds. Notice that, in order to avoid any confounds, decisions related to the sequence of questions are not restricted in this

[6] Notice that, in order to test whether the random allocation to treatments worked, I have used the Wilcoxon-Mann-Whitney test, since it tests specifically for differences in median values of the three samples, without imposing the normality assumption on the distribution of the outcomes.

Table 1.3: Average Score for Baseline Treatment and Fixed Time Treatment groups

	Part 1			Part 2			
	Baseline	Fixed Time		Baseline	Fixed Time		
Outcome	Unconstrained	Unconstrained	Δ_1	Unconstrained	Fixed Time	Δ_2	DID
Score	10.091	9.698	-0.393	10.667	7.147	-3.520	-3.127
Std. Error	0.6854	0.518	0.611	0.684	0.628	0.991	1.164
t			-0.64			-3.550	-2.68
P>t			0.523			0.000***	0.008***

Notes: *Score* is the outcome variable, which sums the points obtained from the correct answers minus the penalty encountered from the wrong answers. Δ_t represents the average difference in the scores, in each part t, faced by the subjects in the groups "Baseline" and "Fixed Time". *DID*: represents the Difference-in-Differences estimate, which is given by $\Delta_t - \Delta_{t-1}$. Legend: On the left, results related to the fist part of the test, $Part_1$. On the right, results related to the second part of the test, $Part_2$. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$

treatment and, therefore, subjects could still choose to switch task. Moreover, if students decide to change question, they will not lose any of the remaining time for the switched task, since the time would be stopped exactly at the moment of the click on the new question.

The aggregate results show, in table 1.3, that imposing a given time per task, in this context, is significantly detrimental to the overall performance. In particular, it reduces the test score by more than 3 points (t statistics = -2.68), which corresponds to the 18% of the average score [7].

As expected from the theoretical predictions, I found a negative and significant effect of the Fixed Time treatment for the overall performance, but small and not significant results have instead emerged for the Fixed Sequence treatment.

In order to enrich the analysis and better characterize the results, similar estimations were made by taking into consideration gender' differences. However, no significant results emerged through performing such sub-group analysis.[8]

1.5.2 Difference-in-Difference-in-Differences

In the previous analysis, I tested whether the predicted average treatment effects were indeed detected in the current experimental context, and I have concluded that imposing a fixed time per task is significantly detrimental, while forcing participants to work sequentially has no effect.

In this part of the analysis, I want to use the information recorded through the mouse-click tracking to better refine the aggregate results. In particular, from the theoretical discussion of the FCFS+FB(θ^*) policy, we have learned that, if the probability of knowing the task is high enough, the individual should complete the task at the first attempt, otherwise she should

[7] Summary statistics are shown in table A.1 in the Appendix A

[8] Results are available upon request.

postpone the answer and return to the task when all the easier questions have been completed. According to this strategy, subjects should not return to each question too often: if they postpone the task, they will return to it when all the easier questions have been completed, and they will continue to work until its hazard will stop to increase; finally, when the hazard of each task has started to decrease and if more time is still available, subjects should then look at those tasks which have received less service, until their completion, and then move progressively to the ones which have received more and more service. If subjects followed this strategy, I should not observe more than three lookups per question, i.e the number of times the subjects has returned to the same question before giving the final answer, especially if she easily knows the answer. Moreover, I expect that students will leave the difficult questions at the end of test, so that the correlation between the time spent on such questions and the time elapsed since the beginning of the test should be positive.

Summary descriptions of the timing and the sequence of each click performed by each subject are described in table A.3 of the Appendix A.

As it is possible to see from table A.3 of the Appendix A, on average students performed more than 3 lookups per question in the first part of the test. This measure gives the number of times students have returned to the same question during the test before giving the final answer. The table also shows information on the average time spent on questions whose answers were missing, wrong and right, respectively. Notice that the time for unanswered questions is the greatest. Finally, on average, students spend about 2 minutes and 20 seconds per question, and the mathematics section seems the one over which subjects generally took most time, while they seemed to find the logic and the verbal comprehension parts easier.

The rest of the analysis continues by using the information detailed above to better characterize the aggregate results.

The first behavioral component that I take into consideration are the lookups. In particular, the predictions from the reference model suggest that students should not return to each question too often, especially if they easily know the answer. Given that I do not know how difficult each question is for each subjects, I look at the aggregate behavior. The mean number of total lookups is near 51, so in order to identify those students who have switched more often, I create a dummy variable that takes value 1 if the total number of lookups exceeds such mean bound. As it is possible to see from table A.4 of the Appendix A, this categorization splits the sample by identifying students who have revised each question more than three times (mean number of lookups larger than 3) from those who have look less often each task. Table A.5 shows that 16% of subjects enters the former group of subjects. In table 1.4, I present the treatments' effects on the probability of giving the correct answer, for those subjects who have looked more than 3 times each question and for those who have not. The table shows that, those students who have looked the questions more frequently (in the first part of the test) have a lower probability of giving a correct answer in both parts ([$Switchers$] and [$Switchers$] x $Part_2$ have negative and

Table 1.4: Treatments' effects on the probability of giving the right answer - Switchers

	(1) $Pr(RightAnswer)$	(2) $\Delta_{2-1}Pr(RightAnswer)$
$Part_2$	.054*	
	(0.031)	
$[Switchers]$ x $Part_2$	-0.163**	
	(0.066)	
$FixSequence$ x $Part_2$	-0.028	
	(0.045)	
$FixTime$ x $Part_2$	-0.208***	
	(0.044)	
$FixSequence$ x $[Switchers]$ x $Part_2$	0.134*	
	(0.080)	
$FixTime$ x $[Switchers]$ x $Part_2$	0.324***	
	(0.084)	
$FixSequence$ x $[Switchers]$		0.219**
		(0.103)
$FixTime$ x $[Switchers]$		0.307***
		(0.108)
$FixSequence$	-0.103	-0.043
	(0.133)	(0.042)
$FixTime$	-0.147	-0.207***
	(0.132)	(0.042)
$[Switchers]$	-0.247*	-0.188***
	(0.145)	(0.072)
$Intercept$	.718***	0.058**
	(0.097)	(0.029)
N	2088	1044
adj. R-sq	0.211	0.170

Notes: *Pr(Right)* is the outcome variable which represents the probability of giving a correct answer during the test. $[Switchers]$ is an indicator variable equal to one if the subject has a mean number of lookups per question greater than three. *Fixed Time* and *Fixed Sequence* are the dummy variable indicating if the subject belong to the respective treatment group. $Part_2$ it is a dummy variable equal to one if the probability in the second part of the test is considered. Standard errors in parentheses. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$.

significant coefficients) as compared to those who switched less frequently between questions (the Intercept and the coefficient of $Part_2$ are 0.718 and 0.054 respectively) when both groups solved the each part of the test without any restrictions. In addition, even if those who had switched less frequently had enhanced their chances of correctly answering the questions in the baseline case, they lose this advantage completely when they are not left free to organize themselves [the coefficients on $FixedSequence$ x $Part_2$ and $FixedTime$ x $Part_2$ are both negative and the latter is also significant and large in magnitude]. On the contrary, if I consider those who have switched more frequently, I see that they reached higher levels of performance both when they are forced to answer sequentially and when they faced a given time per task. In particular, they fill the performance gap that, otherwise, they would have experienced in the unconstrained case with respect to the other group of subjects [both $FixedSequence$ x $[Switchers]$ x $Part_2$ and $FixedTime$ x $[Switchers]$ x $Part_2$ coefficients are positive and significant].[9]

In the second column, instead, results on the change in the probability of giving the right answer are presented for the two groups of subjects. As it is possible to notice, students, who have switched questions more frequently, decrease their chances of correctly answering a question in the second part of the test (-0.188 percentage points) with respect to those who have shown less erratic movement during the first part (intercept + 0.058 percentage points). Moreover, when the latter group is treated, either with a fix time per task or with an imposed sequence, they more than loose what, otherwise, they have gained, especially under the Fix Time scenario [-0.207 percentage points]. Whereas, if we look at the treatments' effects in the former group we find incremental and significant results for both the treatments, exceeding the loss potentially faced in the unconstrained case[10].

Such findings are further confirmed if I look at the final score obtained during both parts of the test. In particular, table 1.5 shows that the Fixed Sequence and Fixed Time treatments are still significant and relevant, if I look at the final score obtained by those students who switched repeatedly across questions, with respect to the case where they are left free to organize their work.

By looking at the table 1.5, I notice that both treatments induce the "frequent switchers" to perform more or less as the other group of subjects, even if we know that, without these interventions, they would have performed significantly worse [the coefficients on $[Switchers]$ and $[Switchers]$ x $Part_2$ are both negative and the latter is also statistically and economically significant].

From the evidence detailed above, it is clear that treatments induce heterogeneous effects, depending on which type of student is considered. In fact, the Fixed Sequence treatment is beneficial for that share of students who have switched and looked up each question more often, suggesting that preventing them from adopting such behavior, in the second part of the test, is significantly beneficial. Concerning the Fixed Time Treatment, I have found that even this schedule could be beneficial for such group of students, although it has not being specifically designed to address such behavioral dimension. In contrast, such constraint is indeed detrimental to those students who better self-organize their answering strategy by not switching repeatedly across questions. A possible reason for this result may be that the imposition of

[9] The F test on the joint significance of $[Switchers]$ x $Part_2$ and $FixSequence$ x $[Switchers]$ x $Part_2$ is not significant ($Prob > F = 0.729$). While the F test on the joint significance of $[Switchers]$ x $Part_2$ and $FixTime$ x $[Switchers]$ x $Part_2$ is significant at the 7% level ($Prob > F = 0.074$).

[10] The F test on the joint significance of $[Switchers]$ x $Part_2$ and $FixSequence$ x $[Switchers]$ x $Part_2$ is not significant ($Prob > F = 0.783$). While the F test on the joint significance of $[Switchers]$ x $Part_2$ and $FixTime$ x $[Switchers]$ x $Part_2$ is significant at the 8% level ($Prob > F = 0.082$).

strict bounds on the time for completing each task may spur students on to be more focused on the task at hand, preventing them from switching repeatedly between questions, which is somehow confirmed, since the mean number of lookups decrease marginally for the Fixed Time group in the second part of the test, see table A.6.

Table 1.5: Treatments' effects on the final score - Lookups

	(1)	(2)
	Score	$\Delta_{2-1}Score$
$Part_2$	1.129	
	(0.189)	
$[Switchers]$ x $Part_2$	-3.329**	
	(1.663)	
$FixSequence$ x $Part_2$	-0.583	
	(1.139)	
$FixTime$ x $Part_2$	-4.112***	
	(1.121)	
$FixSequence$ x $[Switchers]$ x $Part_2$	2.870*	
	(1.809)	
$FixTime$ x $[Switchers]$ x $Part_2$	6.452**	
	(2.123)	
$FixSequence$ x $[Switchers]$		4.764**
		(2.159)
$FixTime$ x $[Switchers]$		6.237***
		(2.260)
$FixSequence$	-0.336	-0.914
	(0.766)	(0.897)
$FixTime$	-0.427	-4.100***
	(0.759)	(0.878)
$[Switchers]$	-1.190	-3.930**
	(0.851)	(1.521)
$Intercept$	10.290***	1.230*
	(0.551)	(0.621)
N	87	87
adj. R-sq	0.158	0.209

Notes: *Score* is the outcome variable, which sums the points obtained from the correct answers minus the penalty encountered from the wrong answers. $[Switchers]$ is an indicator variable equal to one if the subject has a mean number of lookups per question greater than three. *Fixed Time* and *Fixed Sequence* are the dummy variable indicating if the subject belong to the respective treatment group. $Part_2$ it is a dummy variable equal to one if the score in the second part of the test is considered. Standard errors in parentheses. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$.

In the following analysis, I present the results related to the second decision-making dimension in this context, which is the time allocation. Remember that, during the first part of the test, all the students could freely organize their time to answer all the questions. Moreover, as suggested by the optimal policy $FCFS + FB(\theta^*)$, students who want to maximize their performance should intuitively firstly answer those tasks for which they easily know the solution. Equivalently said, they should not spend too much time on difficult questions at the beginning of the test, since this could reduce their available time for answering other questions whose answers can be given more easily.

This reasoning implies that I should observe a positive correlation between the time spent on the question and the time elapsed since the beginning of the test (at the moment when the task is processed for the last time) for the difficult tasks: the higher the time required by the task, the farthest the moment in which she gives the final answer, from the beginning of the test. In order to understand how students allocate time during the test, I constructed measures of time allocation using the information available from the *mouse tracking*. First, I recorded the time each student used to answer each question, by summing up all the seconds spent on each task, even if not consecutively. Afterwards, I measured the time elapsed between the beginning of the test and the last time the subject looked at the question for those tasks whose answers were either wrong or missing, i.e the difficult tasks. As previously underlined, it is reasonable to expect the correlation between these two measures to be positive. This case, indeed, identifies whether students have properly decided not to get stuck on hard questions at the beginning of the test, and, consequently, to save time for answering other questions. In order to identify this group of students, I created a dummy variable that is equal to zero if the correlation between the time spent on the question (whose answer was either wrong or missing) and the time passed since the beginning of the test is negative. The distribution of the students in such an indicator variable is collected in table A.7 of the Appendix A. As it is possible to notice from the table, and according to the above definition, 25% of subjects have a good time allocation, while the remaining share does not allocate time efficiently across tasks.

In the following table, I present the results of the Diff-in-Diff-in-Diff regression estimates for the groups classified, as explained above, on the time-management dimension. In the first column of table 1.6, I see that imposing a given time per task, as a given sequence of work, is detrimental for those students who have shown good time-management skills (the coefficients of $FixedSequence$ x $Part_2$ and $FixedTime$ x $Part_2$ are both negative and significant). Moreover, while fixing a given sequence of work helps in enhancing the performance of those who have bad time-management, the treatment "Fix Time" has no significant effect.[11] The reason why I have not found statistically significant results for the "Fixed Time" treatment could be the fact that, under time pressure subjects may have lower depth of reasoning and thereby the probability of correctly answer the task is reduced. To confirm such intuition, I have looked at how average probabilities of giving right and wrong answers change across groups and across parts. In particular, from table A.8 of the Appendix A, it is possible to confirm that subjects in

[11] The F tests on the joint significance of $[BadTimeManagement]$ x $Part_2$ and $FixSequence$ x $[BadTimeManagement]$ x $Part_2$ and of $[BadTimeManagement]$ x $Part_2$ and $FixTime$ x $[BadTimeManagement]$ x $Part_2$ show that the overall effects are not significant in both cases.

the"Fixed Time" group have actually increased the probability of wrongly answer the task in the second part (when they face a given deadline per task) with respect to the first part (when they are unconstrained). While the other groups do not change their answering strategy, those who faced the pressure of time make more mistakes, confirming the spead/accuracy tradeoff found in the psychological literature [see Maule et al. (2000) for a review].

Table 1.6: Treatments' effects on the probability of giving the right answer - Time-Management

	(1) $Pr(RightAnswer)$	(2) $\Delta_{2-1}Pr(RightAnswer)$
$Part_2$	.110**	
	(0.048)	
$[BadTimeManagement]$ x $Part_2$	-0.118**	
	(0.054)	
$FixSequence$ x $Part_2$	-0.150**	
	(0.071)	
$FixTime$ x $Part_2$	-0.223***	
	(0.064)	
$FixSequence$ x $[BadTimeManagement]$ x $Part_2$	0.192***	
	(0.073)	
$FixTime$ x $[BadTimeManagement]$ x $Part_2$	0.092	
	(0.066)	
$FixSequence$ x $[BadTimeManagement]$		0.145
		(0.094)
$FixTime$ x $[BadTimeManagement]$		0.051
		(0.085)
$FixSequence$	-0.022	-0.115
	(0.030)	(0.083)
$FixTime$	-0.032	-0.194***
	(0.030)	(0.072)
$[BadTimeManagement]$	-0.016	-0.091
	(0.028)	(0.059)
$Intercept$	.674***	0.091*
	(0.029)	(0.049)
N	2088	1044
adj. R-sq	0.0159	0.0159

Notes: *Pr(Right)* is the outcome variable which represents the probability of giving a correct answer during the test. $[BadTimeManagement]$ is an indicator variable equal to one if the subject has shown an inefficient time allocation in the first part of the test. *Fixed Time* and *Fixed Sequence* are the dummy variable indicating if the subject belong to the respective treatment group. $Part_2$ it is a dummy variable equal to one if the probability in the second part of the test is considered. Fixed effects estimation. Standard errors in parentheses. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$.

The findings on the probability of correctly answer the task are further confirmed if I look at the final score obtained during both parts of the test. In particular, table 1.7 shows that students who get stuck on difficult questions at the beginning of the test would have performed much worse in the second part of the test with respect to those who have good time management, if no intervention would have imposed [$BadTimeManagement$x$Part_2$ is negative and significant, while $Part_2$ is of similar magnitude, but positive]. Incremental results of the treatment "Fixed Sequence" may be explained by the fact that, if people know that they can see each question only once, they might become more selective as to which questions they spend their time on. However, further research is necessary to confirm such intuition.

Table 1.7: Treatments' effects on the final score - Time-Management

	(1) $Score$	(2) $\Delta_{2-1}Score$
$Part_2$	2.343*	
	(1.247)	
$[BadTimeManagement]$ x $Part_2$	-2.527*	
	(1.400)	
$FixSequence$ x $Part_2$	-3.243*	
	(1.841)	
$FixTime$ x $Part_2$	-4.231***	
	(1.645)	
$FixSequence$ x $[BadTimeManagement]$ x $Part_2$	4.186**	
	(1.892)	
$FixTime$ x $[BadTimeManagement]$ x $Part_2$	1.609	
	(1.717)	
$FixSequence$ x $[BadTimeManagement]$		3.167
		(2.028)
$FixTime$ x $[BadTimeManagement]$		1.180
		(1.841)
$FixSequence$	-0.334	-2.455
	(0.789)	(1.784)
$FixTime$	-0.390	-3.930**
	(0.777)	(1.554)
$[BadTimeManagement]$	-0.131	-2.115*
	(0.741)	(1.274)
$Intercept$	10.183***	2.055**
	(0.752)	(1.066)
N	87	87
adj. R-sq	0.0159	0.0159

Notes: *Score* is the outcome variable, which sums the points obtained from the correct answers minus the penalty encountered from the wrong answers. $[BadTimeManagement]$ is an indicator variable equal to one if the subject has shown an inefficient time allocation in the first part of the test. *Fixed Time* and *Fixed Sequence* are the dummy variable indicating if the subject belong to the respective treatment group. $Part_2$ it is a dummy variable equal to one if the score in the second part of the test is considered. Standard errors in parentheses. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$.

In the following analysis, I look at the intersection between the two classifications previously described, in order to have a complete typology of the students' behavior. In particular, I split the sample into four categories as follows: students who have good time management and who have a mean number of lookups that is less than the sample mean ("Rational"); students who had allocated time efficiently across tasks but who had switched frequently between questions ("Switchers"); the third group identify those students who have not looked up to questions too frequently but who had poor time management ("Badtime"); finally, the last group is formed of those who have not managed time properly and, at the same time, have looked more frequently each question ("Switchers & Badtime - S&B").

In table 1.8, I present the treatment effects for all the groups of subjects described above, for what regards their final score. The intercept represents the score obtained in the first part of the test by the "Rational" subjects, when they answered the test without any restriction. Notice that this group of subjects obtained a higher score by around 2 points ($Part_2$ coefficient), in the second part of the test, when they were still free to answer the test without constraints. Instead, when such subjects are treated they more than lose what otherwise they will have gained (both the coefficients on $FixedTime$x$Part_2$ and $FixedSequence$x$Part_2$ are negative, large and statistically significant). If I look at other groups of subjects, instead, I notice that they all achieved lower scores than the"Rational", as expected, in the second part of the test, when they could still freely organize their work. In addition, I see that the constraints imposed under the two treatments are indeed incremental for their performance. In particular, the "Switchers & Bad Time" are the ones that are most significantly helped by such restrictions, reaching in the end a performance similar to that of the "Rational" group. From the above evidence, it is clear that taking into account heterogeneity is highly recommended when analysing the effects of such restrictions on workload management. Moreover, the estimates have shown that the average effect will depend on the share of the specific types in the population. Therefore, knowing their distribution is fundamental for inferring which group is driving the aggregate results and for drawing proper conclusions.

1.5.3 Robustness checks

In this section, I will describe some of the robustness checks performed in order to validate the above analysis.

In order to further support the validity of the results, I decided to replicate the above analysis by using a potentially more general classification. In particular, I took advantage of the measures collected in the final questionnaire on specific subjects' characteristics. As previously described, part of the questionnaire was dedicated to the test for impulsive behavior Frederick (2005). We grouped the answers of this test in a unique indicator variable that is equal to 1 if the student was defined as impulsive in all three questions of the test. Summary information on the distribution of students into such categories is shown in table A.9 of the Appendix A.

More than 50% of the students are classified as impulsive in all the three questions of the Frederick's test. In the following paragraph, I will show the results of the Difference-in-Difference-in-Difference regression estimates with such groupings. Table A.10 of the Appendix A summarizes the evidence. Notice that the table shows the average score obtained in both parts of the test for the two groups of students. The students who did not answer impulsively to the test in the questionnaire obtained a significantly higher score during the test, by more than 2 points, suggesting that this type of behavior is extremely relevant in this context.

Moreover, this group of students seems to lose in terms of performance when it is treated. However, if I look at impulsive students, the opposite results are true, with the "Fixed Sequence" treatment having the larger incremental effect.

Such evidence suggests that failing to be impulsive in this context may imply more reflective decisions in the organizational strategy. In particular, as it is possible to notice from table A.11, the total number of lookups is higher for the impulsive students, while their scores obtained in the first part is always lower than the one reached by non-impulsive subjects. Such descriptive evidence might suggest a relation between the impulsivity measure retrieved through the Cognitive Reflection Test and the switching behavior. However, further research should address the validity of this link.

Finally, I have performed a "placebo" analysis by using types of classifications based on variables that should not affect the results, as for example the risk attitude, and indeed I do not find any statistically significant results [12].

To conclude, the results presented in the previous sections are consistent and robust to several checks. Moreover, even adopting a different and, in a way, more general classification, I found significant heterogeneous treatment effects and positive signs for both the imposed schedules, suggesting that they could serve as instruments to correct behavioral mistakes that might emerge in such problems of work-division and attention allocation.

1.6 Conclusions

The results have shown that, on aggregate, imposing a fixed time to solve each task could be significantly detrimental in terms of overall performance with respect to the case of an unconstrained schedule. Whereas, imposing a sequential type of work has no significant impact with respect to the case where subjects could self-organize their work, when tasks are independent.

Contrary to previous studies, however, the present paper has deepen the analysis, refining the above aggregate results by means of the *mouse-click tracking* technique. Especially, the sub-group analyses have shown that the treatment effects are, in fact, heterogeneous and their net impacts change according to individual' types. In particular, the sequential rationale is significantly beneficial to those students who switch repeatedly between tasks and who have shown bad time management skills. The reason is that, by fixing the sequence of work, they have been prevented from switching repeatedly between tasks and, potentially, they have been indirectly helped in better prioritizing the tasks over time. However, both treatments are instead detrimental for those subjects who have shown good prioritizing skills. Given that the

[12] Estimates available upon request

real "Admission Test" constrains subjects to answer each section sequentially, it might be that those students who are good at self-organizing their workload lose from this restriction.

As regards the schedule that fixes a given answering time for each question, I still found positive results especially for those students who switched frequently between tasks and who had shown poor time management skills. Notice that the real assessment test adopted by the University of Bologna fixes a total time for each section. Therefore, since even the Fixed Time treatment is significantly detrimental to the "Rational" students, which are the ones who generally perform better, the university should be concerned by the fact that, at the passing threshold, "Rational" subjects might be penalised by this feature of the design of the test.

Given all the above results, the experiment has shown that imposing a fixed working schedule, either a fixed sequence or a deadline per task, in the workplace may enhance the performance of those workers who lack from efficient organizational skills, while it might reduce the performance of those employees who are efficient in prioritizing and organizing the workload. Therefore, in order to maximize the individual and, consequently, the average output, the employer, should firstly asses the types' composition in her group of workers and then propose individual-specific working schedules, if the composition of the group is rather heterogeneous.

Moreover, the experiment has also shown that the design of the assessment method could impact the types of student who succeeds the selection of the "Admission Test". If the university wants to equalize students' organizational abilities, providing general guidelines before the selection starts might be the solution. However, interesting questions are still open as to the reasons why students do not recognize the optimal answering strategy, and which is the best way to "teach" them.

Finally, using a different, but related, classification based on the standard impulsivity test proposed by Frederick (2005), I found that imposing a given sequence or a given time per task is beneficial for those students who are identified as impulsive, suggesting, therefore, a way to reduce the costs of such behavioral "mistake".

To conclude, such results have put forward new and interesting evidence in term of the heterogeneity of treatments. In particular, future studies that want to extend such research should account for the specific sample composition in order to understand which type is driving the aggregate results and properly to design specific policy for the target group.

Table 1.8: Regressions on Treatments' Groups - Score

	(1) *Score*
Switchers x $Part_2$ x *FixTime*	1.209
	(5.007)
BadTime x $Part_2$ x *FixTime*	0.279
	(1.716)
Switchers&BadTime x $Part_2$ x *FixTime*	8.204***
	(2.624)
Switchers x $Part_2$ x *FixSequence*	3.968
	(5.123)
BadTime x $Part_2$ x *FixSequence*	3.718**
	(1.884)
Switchers&BadTime x $Part_2$ x *FixSequence*	6.015**
	(2.610)
Switchers x $Part_2$	-3.223
	(4.092)
BadTime x $Part_2$	-1.925
	(1.404)
Switchers&BadTime x $Part_2$	-4.956**
	(2.046)
Switchers	-1.307
	(2.076)
BadTime	0.597
	(0.738)
Switchers&BadTime	-1.071
	(1.074)
FixSequence x $Part_2$	-3.230*
	(1.789)
FixTime x $Part_2$	-4.192***
	(1.598)
FixSequence	-0.347
	(0.766)
FixTime	-0.428
	(0.755)
$Part_2$	2.327*
	(1.211)
Intercept	10.200***
	(0.731)
N	87
adj. R-sq	0.210

Notes: *Score* is the outcome variable, which sums the points obtained from the correct answers minus the penalty encountered from the wrong answers. *Intercept* represent the score obtained by the "Rational" subjects of the Baseline group in the first part of the test. *Switchers* is an indicator variable equal to one if the subject has looked more than the predicted optimal level each question in the first part of the test. *BadTime* is a dummy variable equal to one if the subject has shown an inefficient time allocation in the first part of the test. *Switchers&BadTime* represents the intersection between the two above explained categories. *Fixed Time* and *Fixed Sequence* are the dummy variable indicating if the subject belong to the respective treatment group. $Part_2$ it is a dummy variable equal to one if the score in the second part of the test is considered. Standard errors in parentheses. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$.

Chapter 2

On the Allocation of Effort with Multiple Tasks and Piecewise Monotonic Hazard Function

2.1 Introduction

More and more in current everyday life individuals have to manage a series of tasks, whose completion could give direct or indirect reward. For example, students have to solve homeworks, exercises or problem sets for different subjects and they have to dynamically decide how to allocate their effort in order to maximize their grades, without knowing in advance the amount of attention needed to complete each task. A similar problem is faced by workers who have to decide how to manage their daily workload: from which task they want to start to work and how much time they should spend on each task in order to maximize their the final output.

Typically in these decisional problems, the agent faces, at each point in time, the tradeoff of choosing how much attention should be allocated to each task (exploitation) and from which task she should start to work (exploration). Such *exploration - exploitation* dilemma has been usually formalized as a Markow decision process, in which the decision maker takes a sequence of decisions on a process which is evolving over time. The dynamic programming technique provides a tool to solve such type of model [Bellman (1957)]. However, such method suffers the problem of *curse of dimensionally*, since it becomes intractable as the number of states increases.

Reinforcement learning methods have also been used to solve Markov Decision process of this type, however, they have generated solutions that converge only asymptotically to the optimum[1].

[1] For a review see for example Kaelbling et al. (1996).

The present project wants to study if the Multi-armed bandit problem, which is one of the main tool used to captures the exploration-exploitation dilemma, approximates well a realist decision-making scenario. The advantage of using such framework is that it offers a well-known optimal solution, see Gittins and Jones (1979), for different formulations of the same dilemma: like in job scheduling [Gittins et al. (2011)], in labor hiring [Arlotto et al. (2013)] and in judicial case scheduling [Bray et al. (2016)]. In the present paper I look at how individuals decide to allocate their effort through several tasks in a real-word framework, which resembles the "job scheduling" formulation of the Multi-armed bandit problem. Particularly, subjects are observed during a laboratory replication of the "Admission Test" to the University of Bologna, where they have to solve multiple-choices questions on a variety of topics (logic, verbal comprehension and mathematics) in a limited amount of time and they are incentivized to maximize the final sum of rewards.

While such application is completely new in the literature, there are studies which have looked at the behavior of subjects in solving the multi-armed problem using the"arm-pulling" or the"button-pressing" as main task, in the laboratory [2]. The evidence has suggested that heterogeneities across individuals substantially emerge when we look at the exploration-exploitation dilemma and that there are also important and systematic deviations. Interestingly, Horowitz (1973) finds that subjects over-exploit from bad arms while Anderson (2001) shows that individuals experiment less than the optimum.

Studying the exploration - exploitation dilemma using real-word tasks has the advantage to provide evidence on relevant situations which have economic consequences, while it might reduce the similarity with the assumptions and definition of the theoretical framework. However, this paper has tried to balance such tradeoff and it has studied individual decision making in an educationally relevant context by keeping as many as possible of the characteristics of the "job scheduling" formulation of the Multi-armed bandit problem.

Moreover, information on the individual decisions made by subjects is often unobserved or hard to get. This paper has also the advantage to use the electronic version of the "Admission Test" to the University of Bologna and to track the mouse-clicks of each individual, both in time and in space, in order to get such kind of data. This technique has indeed allowed to completely map the solving strategy adopted by each subject and to have interesting behavioral measures on both the exploration and exploitation dimensions. Moreover, by having such information, one could test if the predictions on the rational strategy are indeed identifying highest performance in the sample and to what extent people deviate from such optimum.

The paper is organized as follows. Section 2.2 describes the general "Job scheduling" model and Section 2.2.1 derives the related index policy. Section 2.2.2 characterizes which type of tasks

[2]See for example Horowitz (1973), Meyer and Shi (1995), Anderson (2001), Gans et al. (2007), Steyvers et al. (2009) and Wilson et al. (2014).

are assigned in the experiment and how this relates to the theoretical predictions of the model. In Section 2.3 several empirical tests on the model predictions are performed. Finally Section 3.6 concludes.

2.2 General model

In this section, the general ”Job scheduling” problem is presented, where an agent has to choose how to process a series of "jobs" in order to maximize the total sum of rewards. This problem is typically described as a Markov decision process (MDP).

In more details, the MDP provides a formal framework to model how the decision making of the agent affects the transition of to future states of the outcomes. In particular, at each stage the agent faces a set of possible action she can choose from, with each action affecting differently the transition of the outcome in future stages. The study of MDP starts with the studies by Wald (1950) and by Bellman (1957). For a recent review see Puterman (2014).

The multi-armed bandit problem is a special MDP. In its typical formulation, there is a gambler who has to decide which slot machines to play, for how long she should continue to play from each machine and in which order she should move across the machines. Each bandit gives a random reward from an unknown probability distribution and, at each decision time, the actions the agent could take are to pull or not pull the arm of a given bandit, with the objective to maximize the future stream of rewards. At time zero, the first decision takes place and, after the action is chosen, there is an instantaneous transition to the next stage and an immediate consecutive change in the reward. At decision time t, the agent has to choose which action to pursue from a set of actions $\{0,1\}$. If the action is equal to *0*, the process for that arm is *frozen* at the current state x: no reward is collected and the there is no change in the state of the process. Whereas, if action *1* is chosen, the agent continues to pull that arm. In this case, the arm yields immediate reward of $a^t r(x) = e^{-\delta t} r(x)$, where a with $(0 < a < 1)$ and $\delta (> 0)$ represent respectively the discount factor and the discount rate parameters and the arm transits, after time t, from state x to y with a transition probability of $P(y|x)$.

The optimal solution to the multi-armed bandit process is a "policy" which identifies a decision rule that, for any time t, specifies which type of action must be chosen as a function of the state of the process at time t, in order to maximize the discounted stream of rewards. For each bandit, when the optimal policy suggests to apply action 1 up to time τ and then to apply action 0 thereafter, we say that it defines the optimal *stopping rule* for that bandit and τ is the associated *stopping time*.

Formally, we can express a MDP as follows. The agent has to find the policy rule π, among all possible rules, which maximizes the value of expected future discounted rewards:

$$V(\xi) = \sup_{\pi} E[\sum_{t=0}^{\infty} a^t r_{it}(\xi_{it}(t)|\xi(0) = \xi)] \tag{2.1}$$

The standard way to solve a MDP is through the *dynamic programming equation*[3] which solves the following *dynamic programming equation* by using the *backward induction* technique starting from period T, then $T-1$, and so on:

$$V(\xi) = \max_{\{i=1,\dots,n\}} [r_i(\xi_i) + a \sum_{y \in E_i} P_i(y|\xi_i) V(\xi_1, \dots, \xi_{i-1}, y, \xi_{i+1}, \dots, \xi_n)] \tag{2.2}$$

Unfortunately, finding V() as a solution of the equation 2.2 is not usually possible since there is a *curse of dimensionality problem*: the number of unknowns grows exponentially with the number of bandits n and with the number of the states. However, in Gittins and Jones (1979) the authors proved that an optimal solution to the multi-armed bandit problems is in the form of a *dynamic allocation index*, i.e the *index theorem*. In particular, they showed that, for each bandit, there is a real-valued index, $\nu_i(\xi_i)$ which depends on its current state only. Moreover, they show that, at each decisional point in time, the optimal policy would be to chose to pull the arm of the bandit with the highest index.

2.2.1 Gittins Index and Jobs

In the previous section, we have seen a general formulation of the problem. In this section, specific characterization of the *dynamic allocation index* is provided.

Firstly, consider a *standard bandit process*, denoted by $S(\lambda)$, which is a simplified version of the general problem, in the sense that if the "pull" action is chosen, the arm will pay a reward of λ from that time t onward. Consider the multi-armed bandit process which is formed by only two arms: N_i and the standard bandit $S(\lambda)$. Formally, if the agent chooses to pull the $S(\lambda)$ arm then the payoff would be:

$$\lambda(1 + a + a^2 + \dots) = \frac{\lambda}{1-a} \tag{2.3}$$

Conversely, suppose the agent chooses to pull the arm of the bandit N_i at time 0 and then she has to follow an optimal rule. Consider that if the agent chooses to switch to $S(\lambda)$ at some time τ ($\tau > 0$) later, it must be optimal to do so at time $\tau + 1$ and thereafter. Therefore the maximization problem becomes:

$$\sup_{\tau > 0} E[\sum_{t=0}^{\tau-1} a^t r_i(x_i(t)) + a^\tau \frac{\lambda}{1-a} | x_i(0) = x_i)] \tag{2.4}$$

Specifically, the agent has to choose the stopping time which maximizes the discounted sum of total reward. After some manipulation is possible to see that the value of λ that makes indifferent the agent to pull the arm of either one of the two bandits is the following:

[3] See Bellman (1957)

$$\sup\{\lambda : \sup_{\tau>0} E[\sum_{t=0}^{\tau-1} a^t r_i(x_i(t)) - \lambda | x_i(0) = x_i] \geq 0\} \tag{2.5}$$

The value of λ, which I will denote by $\nu(N_i, x_i)$ represents the greatest price that one would be willing to receive (per period) for not choosing that specific bandit and not pulling the related arm for one or more periods. From 2.5 I have also that:

$$\nu(N_i, x_i) = \sup_{\tau>0} \frac{E[\sum_{t=0}^{\tau-1} a^t r_i(x_i(t)) | x_i(0) = x_i]}{E[\sum_{t=0}^{\tau-1} a^t | x_i(0) = x_i]} \tag{2.6}$$

where τ is optimal stopping time while the numerator represents the expected discounted sum of *rewards* and the denominator indicates the expected discounted *time*, over τ units.

In general, for a family formed by $N_1, ..., N_n$ alternative bandit processes, an *index policy* indicates to pull the arm of the bandit N_k if $\nu_k(x_k) = \max_j[\nu_j(x_j)]$, where x_i indicates the state of the arm and ν_i is the related index.

A special kind of bandit process is the one which gives a payoff at a random completion time with no intermediate rewards. Such bandit is termed a *Job*. In Gittins and Jones (1979) the authors introduced the "Scheduling Problem" and they showed that it could be solved through the *Dynamic allocation index*. The index for jobs is defined by firstly assuming that the agent could switch the job either if completed or if it has received an integer amount of some unit Δ of service time (i.e the amount of attention received so far). We define by $Ve^{-\delta t}$ the payoff the agent receives if the job is completed at time t. Suppose that the job has a processing time distribution F(), with density f() and denote by $\rho(s) = f(s)/[1 - F(s)]$ the *completion rate* of the job, or, similarly, the *hazard rate*. By letting $\Delta \to 0$, namely, by imposing no restriction on the switching times (the so called "preemptive" case), then we have that the index becomes, for $x \geq 0$:

$$\nu(x) = \sup_{t>x} \frac{V \int_x^t f(s) e^{-\delta)s} ds}{\int_x^t [1 - F(s)] e^{-\delta)s} ds} \tag{2.7}$$

Consider the case where the hazard rate is monotonically decreasing in t: $\rho(x) \geq \rho(t)$ for all $(x < t)$ then is possible to show that $\nu(x) = V\rho(x)$, i.e the Gittins index in state x is equal to the immediate payoff in state x and the optimal policy would be not to spend more units of time on the job after x, since its hazard rate is decreasing for $t > x$. Finally, it is possible also to show that $\nu(t)$ is increasing in $t > x$ if the quantity:

$$Q(t, \delta) = \frac{\int_t^\infty f(s) e^{-\delta)s} ds}{\int_t^\infty [1 - F(s)] e^{-\delta)s} ds} \tag{2.8}$$

is increasing in t. Namely, if the jobs have all increasing hazard rates over time and there is no-crossing in all job hazards, the optimal rule is to undertake the jobs in order of decreasing $V_i Q_i(0, \delta))$.

Moreover, Gittins showed that the index rule is still the optimal policy for the problem where each agent has a "queue" of jobs to process and she has to decide how to organize the work in order to minimizes the number of still uncompleted jobs in the queue. Such results, as explained in the next section, was further extended in Aalto et al. (2009) where the authors have proved the optimality for different families of service time distributions.

2.2.2 Piecewise Monotonic Hazard Rate

As we have anticipated in the previous section, in the case of *Jobs*, where the total working time needed for each job is unknown to the agents (who have information just on the time already allocated) the optimal policy would depend on the distribution of the completion time as proved in Gittins and Jones (1979). In particular, if jobs belong to the *Decreasing Hazard Rate - DHR* class, it would be optimal to undertake the jobs in order of least attained service, i.e *Foreground-Background (FB)* policy. Whereas, if the completion time distribution belong to NBUE (New Better than Used in Expectation), as for the case of Increasing Hazard Rate (IHR), the First-Come-First-Served (FCFS) discipline maximizes the number of jobs completed (or equivalently minimizes the number of uncompleted jobs in the queue).

Aalto et al. (2009) extended the Gittins' framework to the class of piecewise monotonic completion time distribution. For the definition of the index in this framework an auxiliary function, $J(x, \Delta)$, the *efficiency function*, with x and $\Delta \geq 0$, must be introduced. Such function is defined by:

$$J(x, \Delta) = \frac{\int_x^{x+\Delta} f(y)dy}{\int_x^{x+\Delta} \overline{F}(y)dy} = \frac{\overline{F}(x) - \overline{F}(x+\Delta)}{\int_x^{x+\Delta} \overline{F}(y)dy} \tag{2.9}$$

with $\overline{F}(s) = 1 - F(s)$ and $J(x, \Delta)$ continuous with respect to x and Δ.

Thus, the efficiency function $J(x, \Delta)$ represents the fraction between (i) the probability that the job will be completed within Δ units of time and (ii) the expected completion time within these Δ units. Then, the Gittins index G(x) is defined by:

$$G(x) = \sup_{\Delta \geq 0} J(x, \Delta) \tag{2.10}$$

with x ≥ 0 and the optimum quantum of processing time is defined by:

$$\Delta^*(x) = \sup\{\Delta \geq 0 | G(x) = J(x, \Delta)\} \tag{2.11}$$

The above expressions define the Gittins index quantum policy, where the optimal strategy constantly picks the job with the highest index and allocates the related optimal quantum $\Delta^*(x)$ of service. Notice that in this case, the policy makes decisions only when the current job is completed or when it expires its optimal quantum of service.

As anticipated before, Aalto et al. (2009) extended the Gittins quantum policy to the case where the service time distributions are non-monotonic and showed how the optimal policy changes for the problem of a single agent facing a queue of jobs . In particular, they considered the case of a piecewise monotonic hazard rate (firstly increasing and then decreasing) and they showed that the optimal policy would be FCFS+FB(θ^*). Specifically, the optimal schedule is to look at the jobs in order of arrival until they receive the optimal θ^* quantum. Those jobs which have received already θ^* units of time are served according to FB, and jobs which have received zero processing time have priority over those which have already been processed for at least θ^* units of time. The proof relies on the fact that we know that if jobs have increasing hazard rate, the agent should continue to process the tasks as long as the hazard continue to increase. After such optimal quantum has been reached and the hazard have started to decrease, jobs are processed in order of least attained service (FB). Notice that θ^* is the quantum Δ of service which maximizes equation 2.11.

Having defined how the optimal policy changes with respect to the hazard rate of the jobs, I now move to the empirical test on the service time distribution and the related completion rate of characterizing the experimental tasks.

2.3 Empirical testing

In the experiment, students have to solve several tasks in order to maximize their performance and decisions on the workload management become crucial for their final payoff. Specifically, individuals have to take a series of decisions which selects which task must be processed first (*exploration*) and how long to stay on each task (*exploitation*) in order to maximize their final payoff. Notice that, not only the maximization problem is resembling to the one of "Job Scheduling" described in Gittins and Jones (1979), but also the characteristics of tasks are quite similar to the ones of "Jobs". In particular, as for jobs, the reward of each task is obtained on completion (which in this case coincides with the provision of an answer) and each task has a random completion time, so that students do not know what is the completion time and they just observe the amount of service received so far. For these reasons, we can use the rule defined by the *Dynamic allocation index* to solve this maximization problem.

In the "Admission Test" case, students proceed by looking at each task in the queue following the order of arrival in the test and they decide which task to undertake by using the information they get at the first look on each question. In the following part of the section we will better characterize such reasoning and an illustrative example, which uses the data from the experiment, will be presented.

In the previous sections, I have underlined how the optimal policy could change according to the characteristics of the service time distribution. In particular, I have noticed that different rules apply depending on the monotonicity of the hazard rate. Thanks to the data gathered through the *mouse-tracking* technique I can see how the probability of giving the right answer is distributed over time in the context of the experiment. In the following picture, the smoothed hazard of the probability of giving the right answer is plotted over time.

Figure 2.1: Hazard Rates by Task

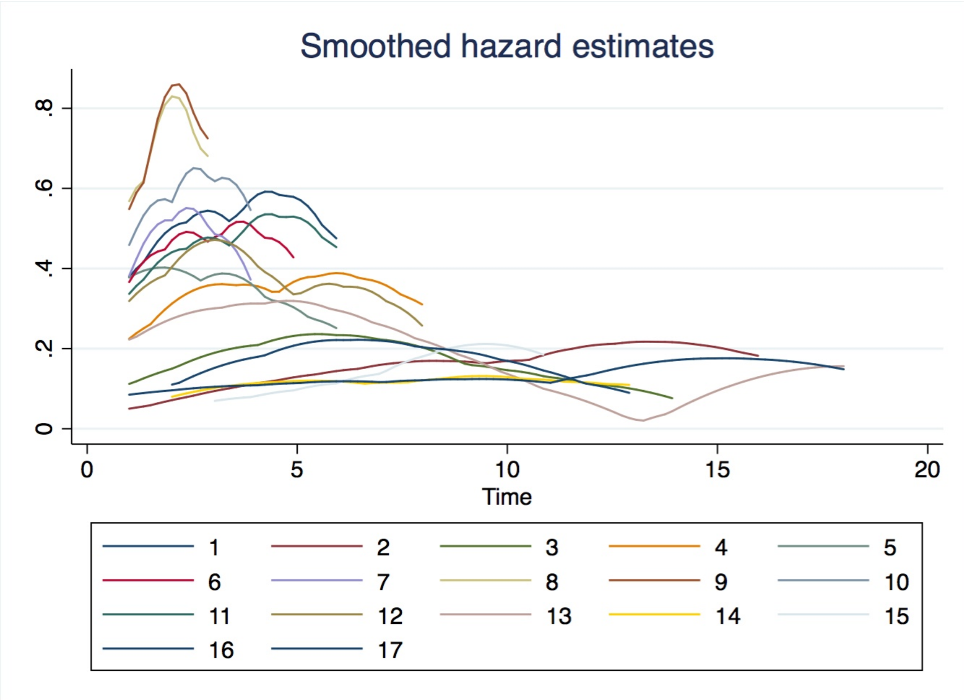

As it emerges from picture 2.1, the hazard of the probability of giving the right answer initially increases and then, after some critical time, it starts to decrease. Notice that, in this graph, the estimated hazards have pooled the observations from all the subjects, so it might be that some task are easy for some subjects but difficult for others. This explains why, for some task there is not a clear pick. However, as it is clearly shown from figure C.1 of the Appendix C, if I consider a subset of subjects a clearer shape of each hazard emerges: firstly increasing and then decreasing[4]. If we think about the context of the "Admission Test", it is intuitive to imagine that the hazard rate for this kind of tasks has such behaviour. In particular, it is plausible that for each task, initially, there is an increasing chance to give the right answer but then, if after this initial time the agent does not know the answer it is unlikely that she will know it afterwards. Moreover, we can think that the critical time would be shorter for easy tasks and longer for difficult tasks but the family distribution would be the same for all tasks.

Such characterization of the *hazard rate* in our context helps in understanding what is the optimal policy for the "Task Scheduling" problem. In particular, the results from Aalto et al. (2009) suggest that, since the completion rate is piecewise monotonic, students should follow a

[4]Figure C.1 is drawn taking just the subjects with an identifier smaller than 4. However similar pictures could be drawn by identifying different sub-groups of subjects. Moreover, notice that for some tasks it was empirically impossible to graph the hazard since some subjects did not provide answers to those tasks.

FCFS+FB(θ^*) policy, where θ^* represents the critical amount of processing time, i.e the optimal quantum Δ^*, that the agent should allocate to each task. Notice that such optimum would eventually change according to the specific hazard rate of each task and for each individual. In this sense, we do not exclude that the hazard rate might be heterogeneous both within subjects and across tasks and across subjects within question. However, since the family distribution of the hazards is the same across tasks and each task has equal independent weight in the accumulation of reward, the Gittins rule, which assigns specific (optimal) quantum of time to each task, will still define what the optimal strategy should be specifically for each subject.

Now I describe how the FCFS+FB(θ^*) policy is applied in such context.

At the beginning of the experiment, subjects do not have any information about the exact completion time of each task. What they only know is that they have to solve 17 tasks in 40 minutes, that each task is valued the same (1 point) if the right answer is provided and the reward of one task is independent from the others. In addition, each question has a multiple-choices format with five available options. With such information set, at time 0 individual i's expectation on the average probability of getting the right answer for task j is equal to 20%, i.e the probability of randomly answering correctly to the question:

$$\begin{aligned} E_0[R_{ij}] &= [R_{ij}|P(Right_{ij}=1)] \times P(Right_{ij}=1) + [R_{ij}|P(Right_{ij}=0)] \times P(Right_{ij}=0) = \\ &= 1 \times 0.20 + 0 \times 0.80 = 0.20 \end{aligned} \tag{2.12}$$

Under this relatively mild assumption, we can now describe how the agent should proceed in the "Task Scheduling" problem.

When she starts the test by opening the first task, she will learn how the hazard rate for that task is distributed. In particular, after the first unit of time spent on the task, she receives information on what is the current probability of correctly answer the task, given that she has spent one unit of time on that task and in that time she did not answer (he learns the *hazard rate*). After this first unit of time has been served, the agent will continue to exploit the same task if the updated probability of completing the task in the next unit of time is higher than the probability of randomly picking one of the answers (without working on task), equal to 0.20 percent[5]. If this is the case, then she will stay on that task as long as the hazard increases, and, she will stop as soon as the hazard starts to decrease. Such optimal length of time for continuing working on the task is indeed defined by the Gittins quantum policy.

Therefore, if the agent chooses to work on the first task, she will continue until either she completes the task or the hazard for that task starts to decrease. After having processed the first task, the individual switches to the next question and the same reasoning applies. Then

[5]This reasoning relies on the fact that, for the tasks we are considering, the marginal increase in probability of giving the right answer is higher for the easy tasks with respect to hard ones, if the same amount of precessing time is spent on such tasks

she switches to the third question and so on, until all the questions in the queue have been viewed. At this point, some of the tasks would be completed (namely those whose answers were given within the optimal amount of time allocated), whereas some others would be still undone, namely those that have been immediately postponed or those which were not completed within the time θ^*. The FCFS+FB(θ^*) policy rule prescribes that, from this moment on, all the remaining tasks must be processed in order of decreasing Gittins index and, for each task, the optimal amount of time, i.e θ^*, prescribed by the index is allocated[6]. The process continues until either the time expires or all the remaining questions have reached the point where their hazards have started to decrease and more time is available. In the former case, the decision process terminates, whereas, in the latter case, the agent should continue to process the tasks in order of increasing service time received, namely from the youngest to the oldest, i.e *Foreground-Background (FB)* policy.

To better understand the policy rule, we will discuss an example by considering the tasks in the test. In the following picture, it's possible to see the smoothed hazard estimates of providing the right answer over time for first four tasks of the Logic section.

Figure 2.2: Hazard Rates in Logic

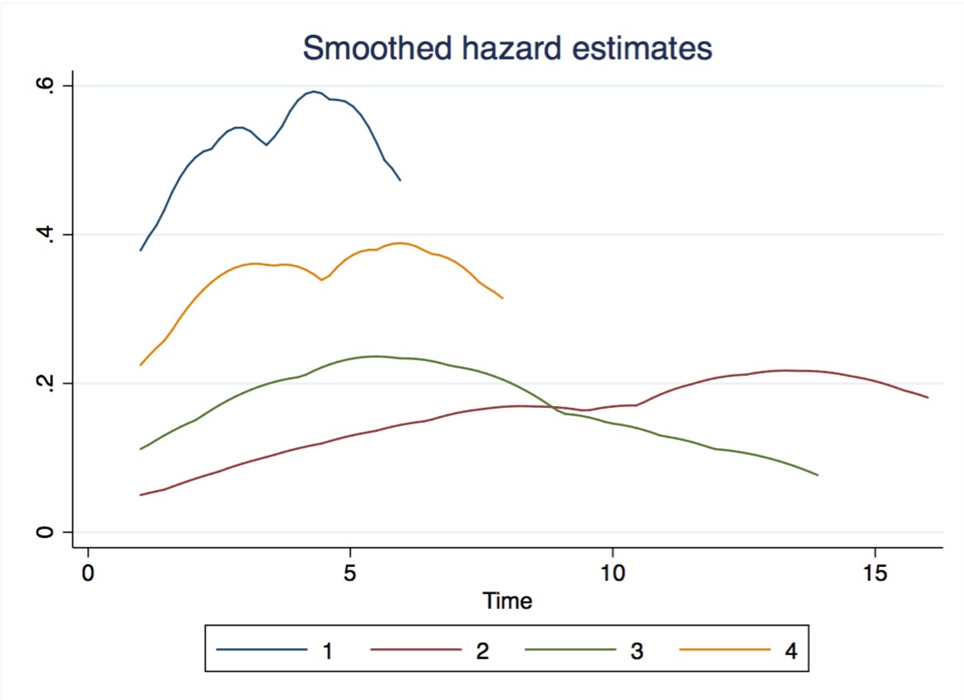

As it is possible to notice, it is likely that the probability of giving the right answer for each task comes from the same family of distribution. What differs across questions is the difficulty, which is defined by the shape of the hazard rate. $Task_1$ and $Task_4$, for example, have very steep

[6]Notice that tasks which have not been analyzed before have priority on those processed by at least θ^* units of time

hazard functions which reach higher maxima in less time with respect to the others.

At time $t = 0$ subjects expect that the probability of giving the right answer is on average equal to 20% for each task after 1 unit of time has been served and, at time $t = 1$, they start answering the test following a First-Come-First-Served discipline. Consider the case of $Task_1$ and $Task_4$ in the example. The hazards of these two questions start to increase almost immediately after the first minutes. Whereas the hazards of $Task_2$ and $Task_3$ are quite flat and they start to moderately grow only after a considerable amount of time has been processed. Suppose that, after the first unit of time spent on the tasks, the agent learns that the hazards of $Task_1$ and $Task_4$ are higher than the average 20% probability, whereas she discover that those of $Task_2$ and $Task_3$ are not. Then the optimal rule implies that $Task_1$ and $Task_4$ should be answered in the first round whereas $Task_2$ and $Task_3$ should be postponed. For simplicity, suppose that the test is composed by just these four questions. At the end of the first round, if the subject has completed $Task_1$ and $Task_4$ she will then continue to work on the task which has the highest Gittins index. Since $Task_3$ has a steeper hazard, its index would be higher than the on of $Task_2$. So she continues to work on $Task_3$ until the hazard reaches the maximum and then, if there is still some time left, she will start to work on $Task_2$ until either the time ends or the task is completed or the task hazard reach its maximum. Notice that, in general, if the agent does not answer within the allocated time, she will stop to work on the task, since from that moment on the hazard will start to decrease, and, in this case, the Gittins policy suggest to provide 0 additional time. Moreover, she will return to such task when all the remaining questions have reached the point where their hazards have started to decrease, and she will process them by following the FB rule (from the youngest to the oldest).

Such case was a simplified example to show, by looking at the average estimated hazard function, how the optimal policy works in this framework. Notice, however, that we do not have individual information on the hazard function of each task. Even though we lack from such data, interesting conclusions about the individual behavior in this framework could still be drawn by using the FCFS+FB(θ^*) as optimal reference.

In the following section, I provide evidence that the FCFS+FB(θ^*) strategy is indeed an optimal policy for this decisional context.

The first prediction that I want to test from the model is about the decision on the *sequence - exploitation* dimension of work. In particular, the optimal policy suggest that individuals should not return to the same question too often: if the probability of knowing the task is high enough she will complete the task possibly at the first look, otherwise she will postpone the answer and she will return on the task when all the easier questions have been completed. The optimality of this prediction is tested by plotting the mean probability of correctly answering the task on the number of times the agent has looked at the task before giving the last answer, i.e *lookups*. The second dimension on which the FCFS+FB(θ^*) policy gives us testable predictions is the *time-allocation - exploitation*. The optimal policy predicts that students should proceed

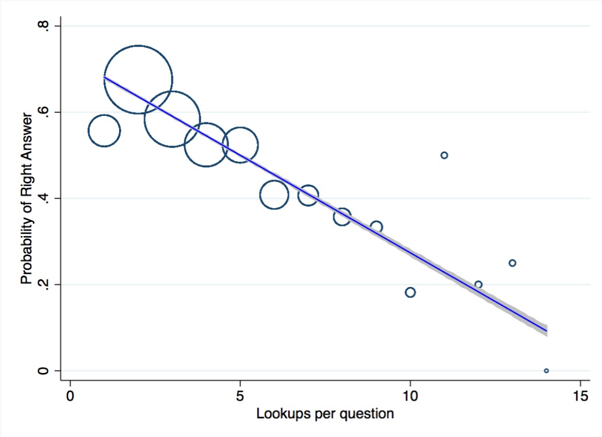

Figure 2.3: Right Answer and Lookups

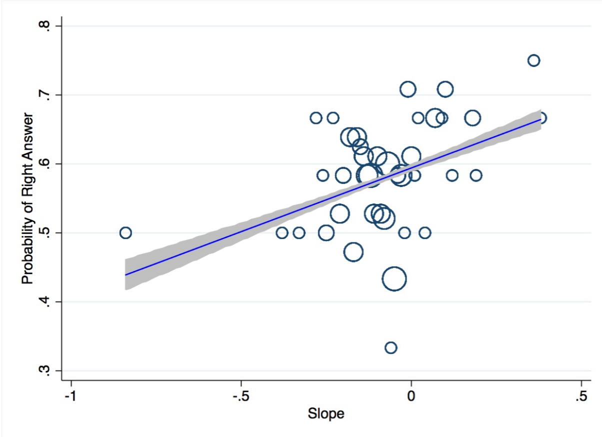

Figure 2.4: Right Answer and Time-Management

by answering those tasks that have a higher probability to be correctly completed (the easy tasks) and by postponing the difficult ones. Such prediction is tested by first computing the correlation between the time spent on the hard tasks and the time passed by from the beginning of the test[7] and then by plotting the mean probability of correctly answer the task for each level of the correlation. If students allocate time properly we should find a positive correlation between the two: students the probability that a question is correctly answered increases with the postponement of the hard tasks, namely those whose completion requires more work, later in time.

Figure 2.3 and 2.4 test these predictions.

As it is possible to see from the above pictures, there is a sharp negative relationship between the mean number of lookups and the probability of giving the right answer. Moreover, the majority of subjects have looked a maximum of 3 times each question, which is the predicted maximum value from the theory. From figure 2.4 it is possible to notice instead a positive relation between the probability of correctly answer the tasks and the postponement of the difficult questions. The fact that we observe a lower probability of a correct answer for tasks which has been switched more often and for subjects who have not given priority to the easy tasks are indeed suggestions of the optimality of the FCFS+FB(θ^*) policy. This result is further confirmed in the following picture where I consider the relationship between the final score obtained by subjects who have looked a maximum of three times each question and who have postponed the difficult tasks at the end of their time.

As it is possible to see from picture 2.5, there is substantial difference in the score obtained between those students who have looked more than three times each question and those who have not. Figure 2.6 shows instead the difference in the final score obtained by those who have

[7]To compute such correlation I use the mouse-click data. In particular, I calculate the relation between the time spent on the hard tasks (whose answers where either missing or wrong) and the time passed by from the beginning of the test running simple OLS regressions and forcing the slope through the origin.

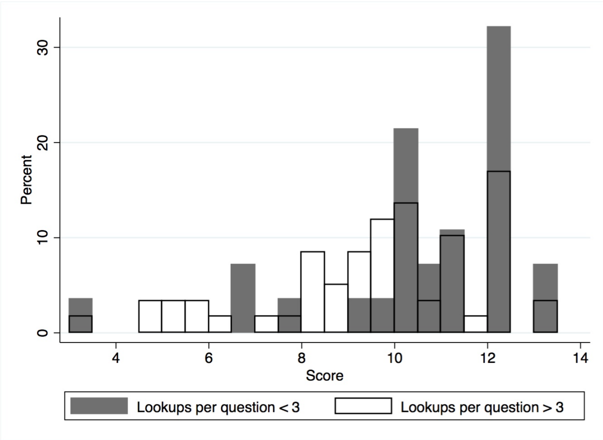

Figure 2.5: Score and Looklups

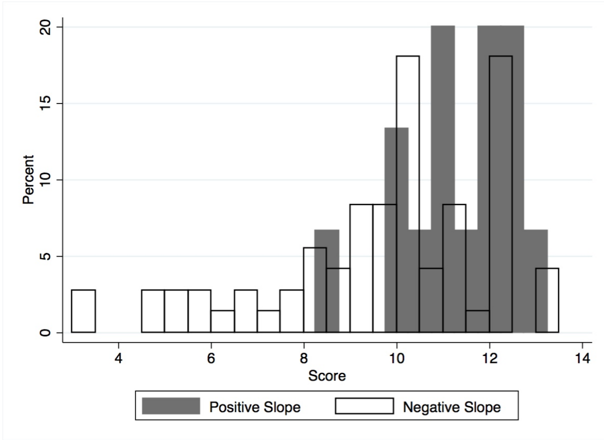

Figure 2.6: Score and Time-Management

given priority to the easy tasks with respect to those who got stuck on difficult questions at the beginning of the test. Such result further confirms that the FCFS+FB(θ^*) policy provides an optimal solution to the optimization of the *dynamic* stream of rewards.

From the above tests we have learned that students who seem to have followed the FCFS+FB(θ^*) policy regarding both the choices related to the sequence of work and the time allocation across tasks, have a higher probability of giving the right answer on each task (and a higher score), with respect to students who have followed different strategies. In the analysis that follows we proceed by focusing the attention on the deviations from such optimal behavior in order to better understand what have characterized the students' decisions.

In order to better understand how students behaved during the test, I constructed the contour plots of a subgroup of tasks, namely those of the Logic section, showing the number of lookups and the length of time spent on each task, over time. The first graph, shown in picture 2.7, plots how the lookups are distributed over time, for correctly and incorrectly answered tasks separately.

It seams that, even if students return to answer the hard questions mostly later in time, (the darker areas are later in time, both for wrong and right answers) there is still a share of students who has returned soon on these tasks and this is even more accentuated for the "Wrong" panel. Such evidence possibly points out why we have found negative relationship between the number of lookups and the score. In particular, we can think that students who have looked more frequently to the tasks are likely to do so more for the difficult questions and at the beginning of the test and such behaviour may have caused them to loose in terms of overall performance. Such intuitive result is further confirmed if we look at the time-allocation: not only they returned frequently to the difficult questions but they also spent a considerable amount of time on such questions still at the beginning of the test. The countour plot below, in fact, shows that even if students on average postponed some of the difficult questions, there is a share of them that "got stucked" on such questions at the beginning of the test, as it is possible

Figure 2.7: Contour plot Lookups in Logic

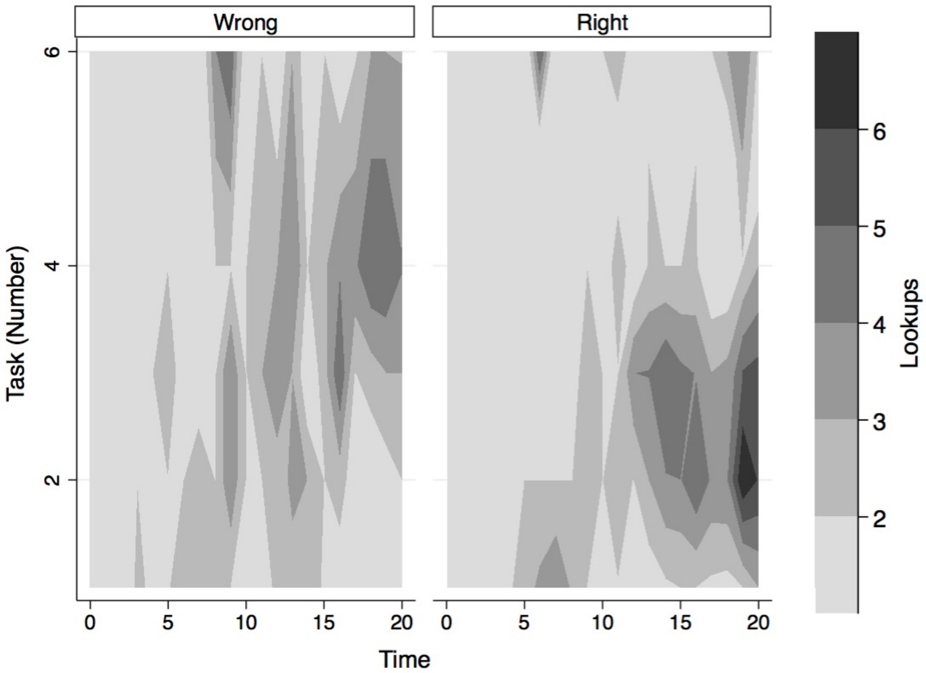

to see from the left panel of figure 2.8.

Such evidence has suggested that on average the predictions of the Multi-armed Bandit model for the "Task Scheduling" with piecewise monotone hazards rate are useful to identify best performance. However, interesting deviations arise: a share of subjects tends to re-process the difficult tasks at the beginning of the test, spending quite a large amount of time. This evidence is in line with previous finding in the literature which has shown that subjects over-exploit from bad arms [see Horowitz (1973)].

2.4 Conclusion

The present chapter examines how the Multi-armed Bandit formulation for the "Job Scheduling" problem could be used for understanding optimal behavior in a real-case scenario. In particular, I firstly describe the general formulation of the standard *scheduling* problem and the related derivation of the optimal policy through the Gittins Index, see Gittins et al. (2011). Then, I discussed how such index policy changes in accordance with the hazard rate of the tasks, using the results from Aalto et al. (2009).

In the empirical section, I firstly provide evidence on how the task in the "Admission Test" to the University of Bologna, resembles the "Jobs" of the standard *scheduling* formulation of the Gittins Index and then I discussed how the optimal policy applies in the experimental context. By using the mouse-click data, I was able to create interesting measure of organizational behavior in the experiment and to test if specific predictions of the model identify optimal outcomes.

Figure 2.8: Contour plot Time on task for Logic

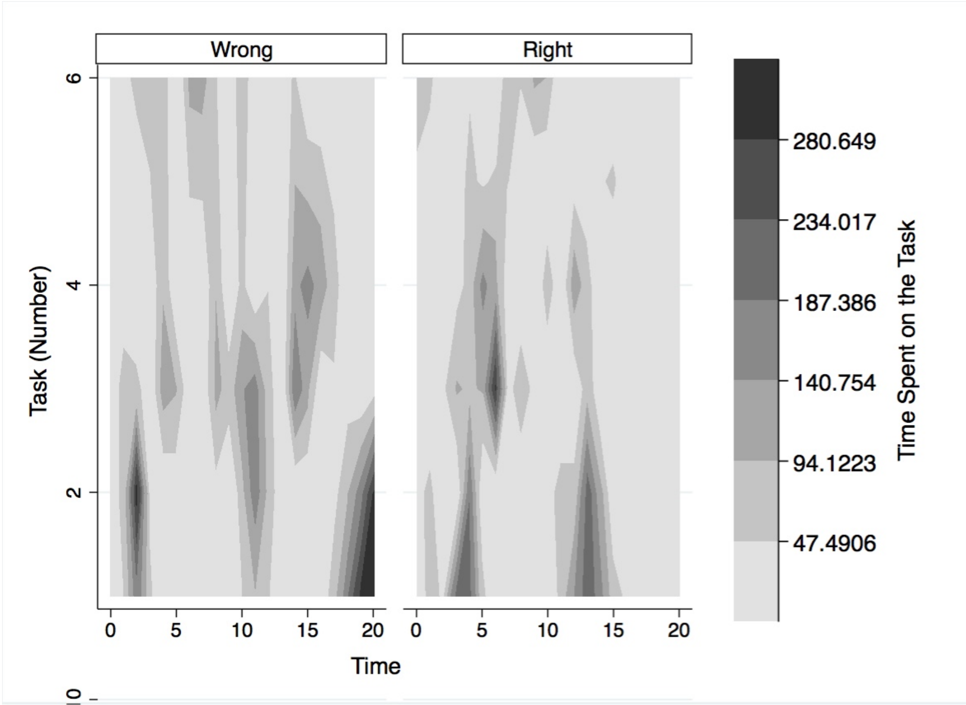

In particular, by characterizing the *exploration* and *exploitation* dimensions with behavioral measures, such as the number of lookups per task and the correlation between the time spent on the task and time passed by from the beginning of test for the "difficult" tasks, several interesting results are observed. Specifically, the data shows that the probability of correctly answer a task is negatively correlated with the number of times the subject has returned to that task and it is positively related with the postponement of the hard tasks.

In addition, I found that such high number of lookups and the inefficient time-allocation are principally in place for the difficult tasks and somehow at the beginning of the test, which, in turns could provide a possible explanation of why costs in terms of performance are encountered.

To conclude, I have tested the prediction that individuals' performances are higher under the FCFS+FB(θ^*) policy and I have found results in line with the optimality of such rule. In addition, I have discussed how students deviate from such policy and why such deviations may have generated real costs in terms of overall output.

In Chapter 1, I use the results from this chapter to understand how restrictions on the individual decision space, which theoretically should negatively affect the performance of subjects, are indeed empirically related to the final outcome.

Chapter 3

The Effects of Scholarships and Tuition Fees Discounts on Students' Performances: Which Monetary Incentive works Better?

3.1 Introduction

Higher education is often promoted to foster the personal development and to increase the wealth of a society more generally. Given the extensive private and social benefits resulting from the human capital investment, governments in many advanced countries have launched several programs aimed at opening the access to the educational system and at fostering the success of students into the programs.

The Human Capital model [see Becker (1975)] suggests that choices of investment in education are essentially driven by a benefit and cost analysis, where the individual compares the direct (tuition fees, equipment, accommodation, etc.) and indirect (forgone labour earnings) costs with the expected benefits of such investment.

Broadly, governments have the general aim of promoting financial aid programs in education in order to decrease the cost of college, especially for liquidity-constrained individuals and, consequently, to induce additional students to enrol into and complete the educational program.

Studies on how monetary incentives are more or less likely to alter behaviour (in the desired directions) are now prominent in the literature [see for example Gneezy et al. (2011)], especially for educational policies. However, while the evidence on the effects of such incentives on *enrolment* has reached a positive consensus [see Dynarski (1999) or Bettinger (2007)], the empirical literature regarding the effects on performance has not reached this point so far. Indeed, the secondary goal of such incentives, especially if minimum requirements for non-returning the benefits are set, is to decrease the need of working during college and, therefore, to increase the hours of study with respect to those of work, inducing, consequently, more studying effort (and possibly better performance). However, results on post-secondary education are still divided [for a review see Fryer et al. (2011) and Dynarski and Scott-Clayton (2013)].

The scope of this paper is to analyze the effects of different educational policies on the performances of similar students, at the end of the first year of their tertiary education. In particular, I exploit the eligibility conditions for different types of scholarships and for several levels of tuition fees discounts, by applying a Regression discontinuity Design and to estimate the causal effects of such different treatments on student performances in a clean identification environment.

The types of institutions analysed in this research are scholarships and tuition fees discounts that are public funded and provided with the purpose of responding to the "Right to Study" principle recognized in the Article 26 of the Universal Declaration of Human Rights, undersigned by all the members of the United Nations: [1]

> *"Everyone has the right to education. Education shall be free, at least in the elementary and fundamental stages. Elementary education shall be compulsory. Technical and professional education shall be made generally available and higher education shall be equally accessible to all on the basis of merit. Education shall be directed to the full development of the human personality and to the strengthening of respect for human rights and fundamental freedoms. It shall promote understanding, tolerance and friendship among all nations, racial or religious groups, and shall further the activities of the United Nations for the maintenance of peace. Parents have a prior right to choose the kind of education that shall be given to their children."*

The advantage of studying this particular types of interventions is firstly to have income-based programs, while most of the previous studies have focused on merit-based financial aids. In particular, I observe the performance of students positioned in a wide range of the Italian social ladder, approximately from 0 €to 40.000 €annual income, without limiting the analyses to the subjects in the top quantiles of the ability distribution. Furthermore, the second advantage is that such policies have a large coverage since all the public funded Universities in Italy should provide these means of financial supports to students and, usually, the share of students who benefit from these programs is quite sizeable [2]. Therefore evaluating such implemented programs could give an important contribution for the policymaker, especially given the large margin of impact.

[1] http://www.un.org/en/universal-declaration-human-rights/index.html

[2] At the University of Bologna, in the academic year 2009/2010 was equal to 13,3%. Elaboration of data from the Italian Ministry of Education Research and University: http://www.ossreg.piemonte.it/doc_02_02_02.asp?nid=7

The study by Mealli and Rampichini (2012), analyses data regarding four Italian universities and they show that the "Right to Study Scholarship" prevents students from dropping out, a result which was further confirmed by Sneyers et al. (2016) through a propensity score matching approach. Moreover, the latter study has also shown a positive and statistically significant effect of the grant on the performance of the recipients with respect to the non-eligible students. The present study wants to add new evidence on the academic performance using a clean identification environment and also by looking at a different incentives scheme, namely tuition fees discounts.

For what concern the literature on tuition fees discounts, at best of my knowledge, the only paper that empirically evaluates a similar type of intervention in a quasi-experimental set-up at tertiary education is the one by Garibaldi et al. (2012), which studies the effects of an increase in college cost, in response to delayed graduation, on on-time completion rate at private University in Italy. The authors find that the students who are enrolled in their final year of study and who may potentially pay a higher fee for an additional extra year of education, are more incentivized to finish on-time, a result which is line with this study.

The results show that increasing the scholarships to low-income students does not affect performance. Such result is mainly explained by the fact that the performance of such groups of students is centered near the necessary requirements for non-returning the benefit. This suggests that students put just the minimum amount of effort required to keep the benefit, independently from the level of the aid received. In addition, I find that by providing no scholarship and letting the students pay half of the tuition fees per year, increases the effort put by these students, in comparison to the level put by those who receive the low scholarship and do not pay any tuition fees. Such result might be intuitively explained by the fact that students who are paying some of the tuition fees might exert more effort to finish on time, since they do not want to suffer the cost of one additional extra year of education, with respect to those who are actually not paying for their education.

The paper is organized as follow. Section 3.2 analyses the institutional framework and Section 3.3 describes the data and the methodology. In section 3.4, I will present and discuss results on the causal effects of monetary incentives on different measures of academic achievement. Section 3.5 discusses the findings and Section 3.6 concludes.

3.2 The Institutional Framework

Tertiary education in Italy is generally accessible to students with a High-School diploma, independently of the type obtained (scientific, classical, professional). The Italian Constitution acknowledge, in art. 34, the "Right to Study" principle of the Universal Declaration of Human Rights by guaranteeing the access even to students at the bottom of the social ladder and reports that:

> "Pupils of ability and merit, even if lacking of the financial resources, have the right to attain the highest grades of studies. The republic furthers the realization of this right by providing scholarships, allowances to families, and other means, to be assigned through competitive examinations."

To meet the purposes of providing equal opportunity and fair access, the Emilia-Romagna region, which is the region where the University of Bologna is located, offers different types of services: allowances for international mobility, housing and meals services, services for people with disability, vouchers for education programs (Master, High-level education, etc.), fiduciary loans, and part-time working possibilities. The present project focuses on scholarships and tuition fee discounts offered for attending the University of Bologna. In particular, the region offers such benefits, with different degrees of generosity, for those students who want to enter a degree program in this University centre, conditioning the awarding only on the income level of the applicants.

The "Right to Study" service in Emilia-Romagna has generally a large coverage, both in terms of the number of recipients and of the amount of financial resources provided. For example, in the academic year 2008/2009, 13.475 students received a scholarship over a total of 77.892 students in the University centre (17,3%). [3] Furthermore, to finance just the scholarship program for the same academic year, the Region of Emilia-Romagna collected: from national resources €151.986.000, from regional resources €158.120.201 , from regional taxes €171.085.441, making a total of €481.191.642. [4]

In this perspective, it becomes quite important to know how these public financed benefits shape students' incentives and, in particular, whether they have any effects on students' academic achievements and which level or which kind of instrument is more effective in shaping the academic performance. This question is of paramount importance in Italy, given that the students' profiles and performance are quite peculiar in this country. In fact, it has been estimated [5] that, at the Bachelor level, 42% of students are "Fuori Corso", i.e those who are enrolled in the University system beyond the legal length of degree program, and that the average time to complete the Bachelor degree is of 5.1 years instead of 3 and the Master degree is of 2.8 years instead of 2. In addition 21,7% of students drop-out from the system. In this context, therefore, properly designed public policy could play a role not only in providing fair access to the University but also in fostering students' performances and ultimately in enriching the national human capital, as through the prevention of university drop-out and the reduction of the share of "Fuori Corso" students.

The regional agency appointed for the distribution of the "Right to Study" services is ER.GO

[3] Data from "Ministry of Education, University and Research - MIUR"
[4] Data from "Ministry of Education, University and Research - MIUR"
[5] AlmaLaurea - Annual Report on University' Graduates 2013

and from 2008 onward, the agency has fully covered all the scholarships and tuition fee discounts'applicants, a 100% successful rate. The application for all the types of the benefits should be made before the academic year starts. [6] Results on allowances eligibility are published within few months. For what concern the scholarships, the first instalment (50% of the yearly allowance) is paid due the end of the calendar year, while the second half of the financial transfer is bind to the satisfaction of precise credit requirements, which are known *ex-ante*, and it is eventually received before the start of the successive year of study. The same credit requirements apply to those students who are receiving a tuition fee discount.

In the following paragraphs, I describe the structure of the benefits offered by ER.GO, starting from the design of the scholarships incentive scheme.

The scholarships vary according to the student "status", which is defined depending on the declared residence: "In sede", which identifies the students who live in the city where the University centre is located or who do not live more than 45 minutes far away from the University centre (by public transport); "Fuori sede", i.e the students who live more than 90 minutes far away from the University centre (by public transport), "Pendolari", i.e the students who are commuting and in particular who are living from 45 to 90 minutes away from the University centre (by public transport). Within groups, "In Sede", "Fuori Sede", "Pendolari", scholarships are assigned according to three thresholds on one income indicator of the family, ISEE. Furthermore, eligibility is always conditional on a maximum value of a further wealth indicator of the family, ISPE, which should not exceed €40.000,00. The ISEE indicator is given by the annual after-tax income plus the 20% of family assets and it is adjusted for the family size by means of an equivalence scale; the ISPE is instead an indicator based just on the family assets (movable and real properties) and it is also adjusted for the family size by means of the same equivalence scale. The amount of scholarships, specified by ISEE thresholds and student "status", are summarized in table 3.1:

Table 3.1: Scholarships' Assignment

ISEE Thresholds	"Fuori Sede"	"Pendolari"	"In Sede"
Up to €12713.21	€5073.78	€3043.88	€2255.11
From €12713.21 to €15386.29	€3942.83	€2420.89	€1828.83
From €15386.29 to €19152.97	€2811.88	€1796.93	€1402.53

The system provides the students who have an ISEE indicator below €19152.97 and who are receiving a scholarship, a full discount, i.e 100% discount, on the yearly tuition fees. Students with an ISEE indicator above €19152.97 and with an ISPE indicator of maximum €60000.00,

[6] At the University of Bologna all the courses last from September/ October to June/July of each year.

can instead apply just for a tuition fees discount, according to the ISEE threshold described in table 3.2

Table 3.2: Assignment of Tuition Fees Discount - 2009

ISEE Thresholds	Fess Discount
From €19152.98 to €22500	50%
From €22501 to €26000	40%
From €26001 to €30000	30%
From €30001 to €35000	20%
From €35001 to €40000	10%

Figure D.1 of the Appendix presents both the benefit schedules for illustrative purposes. The figure shows how the scholarship levels and the tuition fee discounts change as a (discontinuous) function of the income indicator, ISEE. The figure considers as benchmark case a tuition fee equal to €3000, which is approximately near to the average full tuition fee for attending the first year of tertiary education in Italy[7]. Firstly, consider that the sample of scholarship recipients, namely those with an ISEE smaller than €19152.97 (solid blue vertical line), are exempted from paying any tuition fees, so they receive a benefit of +€3000, i.e a 100% fee discount, on top of the scholarship amount awarded, which depends on the income indicator and on the student "status". By considering the schedule of the scholarship recipients, up to the first threshold, the average scholarship is of around €3943, while immediately after that threshold the average scholarship decreases to €2421[8] . This feature gives rise to the first discontinuity on the relationship between the benefit received and the income indicator ISEE. The second discontinuity arises from the fact that, after the second threshold, the scholarship is lowered from €2421 to an average of €1839, thus generating a difference of around €600. The highest discontinuity in the benefit received is at the third threshold, where students with an income indicator around €19152.97 receive extremely unequal benefits. Indeed, the difference is on average of around €3340 [9]. Finally, after such threshold, the discontinuity in the benefit around all the subsequent thresholds is of about €300.[10]

[7]Estimate from the *7th Report on the costs of the Italian universities* of the national non-profit organization "Federconsumatori"

[8]There are no change instead on tuition fee discount, which, for the scholarship recipients is always equal to 100%

[9]Given by the average of the lowest scholarship €1839 and the 50% percent difference in the tuition fee discount, namely around €1500

[10]The schedule of the tuition fee discount slightly increased from the academic year 2009/2010 to 2010/2011, see table D.2 of the Appendix. However, such change do not represent a problem for the estimation since in my sample I do not have observations which enter the small extra ISEE windows.

These discontinuities are the source of exogenous variation which I exploit in identifying the effect of changing the benefit level, even within the same type of program, on the academic performance of students in the first year of their studies. The academic performance is observed both in term quantity and quality. In particular, data on credits and on the weighted average GPA at the end of the first year of study are used to assess the effects of the benefits on the speed and on the accuracy of any change in performance.

All the recipients, regardless of the type and the level of the benefit received, are subjected to given performance requirements for non-returning the monetary aid, after the first year of their studies. In particular, for the undergraduate freshman, scholarships and tuition fee discounts are assigned just on the basis of the two economic indicators, ISEE and ISPE. The master freshman should have obtained, in addition, at least 150 credits in the undergraduate program by the time of the application. After the first year, the benefit maintenance is conditioned to precise yearly credit requirements. Namely, undergraduate students should obtain 25 credits (out of 60) by the end of the first academic year and master students should obtain 30 credits (out of 60) by the end of the first academic year, independently from the types and the levels of the benefit received. Therefore for all the second year applicants (new and first year recipients), the assignment of the benefits is conducted on the basis of the two economic indicators, ISEE and ISPE, and on the basis of the credits obtained in the first year of study. Even at the second year, recipients are required to obtain a target number of credits, namely 80 credits (both for undergraduate and master program), in order not to pay back the benefit. Students enrolled in the third year of the undergraduate program should obtain at the end of the academic year 135 credits not to refund ER.GO of the amount of the benefit received. In table 3.3 and table 3.4, I describe the load and the requirements of credits by year of study and by type of course.

Table 3.3: Credits load by years and programs

Type of course	1st Year	2nd Year	3rd Year
Bachelor	60	60	60
Master	60	60	-

Table 3.4: Credit requirements by years and programs

Type of course	1st Year	2st Year	3rd Year
Undergraduate	25	80	135
Master	30	80	-

3.3 Data and Methodology

The data are provided by the regional agency ER.GO which is in charge of delivering the "Right to Study" benefits, namely scholarship and tuition fee discounts, and by the University of Bologna. These are administrative dataset containing detailed information on the characteristics of each applicants such as the income indicators, ISEE and ISPE[11], the academic performance, both in term of quantity (credits) and quality (GPA), students' demographic characteristics and university-related variables: high-school grade, the macro-region of origin (north, centre, south, Islands), fields of studies, levels of degrees, type of scholarships obtained ("In Sede", "Fuori Sede", "Pendolari"), type of tuition fee discount received ("50%","40%", etc).

The data covers the academic years from 2009/2010 to 2010/2011 and the sample includes all the students admitted in all the twenty-three faculties of the University of Bologna, following their career during such timespan. [12]

Table D.1 in the appendix provides descriptive statistics for the sample of students in the first year of their studies, *per* threshold and *per* type of benefit received. On average 60% percent of the students are female, of approximately 21 years old. The geographical distribution is quite asymmetric as 50% of students are coming from the North, around 20% from the Center, 20% from the South of Italy and 6% form the Island (Sicily and Sardinia). The High School grade in Italy ranges from 60 to 100; the sample mean is around 80 points.

The average number of credits obtained during the first year of enrolment is around 27 in the full sample and the weighted GPA is of about 26[13].

In order to study the effect of scholarships on students' performances, I use a Regression Discontinuity Design. The RDD has been largely used in economics and behavioral sciences and was firstly introduced by Thistlethwaite and Campbell (1960). The attractiveness of such design is that it allows to identify and estimate treatment effects in a context similar to a formal randomized experiment. The identification and estimation issues for the RDD were formalized in the work of Hahn et al. (2001), which describes the minimum set of conditions under which it possible to non-parametrically identify the treatment effect.

The idea of an RDD is to exploit discontinuities in the relationship between an assignment variable and a treatment variable. In the present context, these are respectively the income indicator, ISEE and the level of the benefit received. The intuition is that if the treatment (scholarships or tuition fee discounts) is expected to have an effect on given outcomes, as for credits and GPA, there should also be a discontinuous relationship between the outcomes and the assignment variable, ISEE.

[11]The information on income is subject to legal verification from the agency and the calculus of the ISEE and ISPE indicators must be certified by a professional institution.

[12]From September 2012, the University of Bologna has changed the organisation from 23 faculties to 11 Schools and 33 Departments.

[13]Grades in the Italian university system range from a minimum of 18 to a maximum of 31

In order to test for discontinuities and to identify the Local Average Treatment Effect (LATE) in the data, I apply a local regression, using different polynomial forms, following the parametric model of the form:

$$Z_i = \alpha + \gamma_1 D_i + \gamma_2 F(Y_i - c) + \epsilon_i \text{ where } |Y_i - c| \leq h \tag{3.1}$$

where the set of outcome variables, Z_i, consists of number of credits and GPA obtained in the first year of enrolment, Y_i is the ISEE centred around the threshold c, $F(\cdot)$ is the polynomial order of the regression [14], h is the bandwidth used and D_i is a dummy variable taking the value 1 for $(Y_i - c) \leq 0$. Under this specification, it is possible to demonstrate that:

$$\gamma_1 = E[Z_i|Y_i - c = 0^+] - E[Z_i|Y_i - c = 0^-]$$

identify the level change in the outcome variable at the discontinuity and it is an unbiased estimator of the LATE.

3.4 The Effect of the Benefit Level on Performance

3.4.1 Graphical Evidence

The first graphical evidence on how performances are affected by the benefit level is drawn by picturing how credits earned and GPA are related to ISEE in each academic year. Figures D.2-D.4 show how the data on credits are related to ISEE for 2009 and 2010 separately and jointly. From these graphs, it is possible to notice that the amount of credits obtained by students at the end of their first year of study is relatively flat for the scholarship recipients and basically centred around 25-30. Even if the level of the benefit received decreases at the first two thresholds, the graphical evidence shows that there is no a sharp change in the credits earned by students. Notice that the mean performance of these groups of students is actually close to the first year credit requirements for non-returning the benefit, see table 3.4. On the contrary, at the threshold where there is the highest inequality in the benefit received, we notice a high and somehow persistent jump in the performance of students. In particular, students from the right-hand side of the third threshold earn consistently more credits than those on the left-hand side, even if they are receiving a significantly smaller benefit. Regarding the subsequent thresholds, the graphs are somewhat more noisy but the amount of credits earned seems to be constantly centred around 40, suggesting no sharp effect in decreasing the discount on tuition fee.

[14] Several functional forms are tested to determine which specification better fit the data. Results are shown in the Appendix and will be discussed in the following sections

To investigate whether the results on credits may have generated any side effects on the quality dimension of the study effort, similar graphs are drawn for the GPA outcome. Figures D.5-D.7 summarize the evidence on GPA. As it possible to notice, in such case data seems to be more noisy, especially at the top of the income indicator, and it is not possible to catch clear discontinuities, as for the case of credits. Notice, however, that there seems to be a growing tendency in the GPA with respect to the ISEE since those on the right tail of the income distribution have a higher mean GPA with respect to those at the bottom[15].

It is worth noting that the evidence presented in these figures is not specific to a particular year (both for credits and for GPA) since similar patterns appear in both years, even if the pool of recipients has changed. However, a formal regression analysis is needed to have clearer evidence, especially since there are some noisy patterns around certain thresholds, and the next subsection will present such proper analysis pooling all the yearly data. Figures D.8 and D.9 include the local linear regression of the underlying individual observations with a triangular kernel and optimal bandwidth selection from Calonico et al. (2014b).

In order to explore whether such results are driven by any discontinuities in the ex-ante characteristics of the students, several graphs on the distribution of the numerous covariates with respect to ISEE are reported. Indeed, one essential advantage of using a Regression Discontinuity approach is that the main assumptions, needed to correctly identify the parameters of interest, have direct testable implications. The first prediction is that the conditional expectation of all the ex-ante determined characteristics must be smooth around the thresholds. Secondly, there should be no discontinuity in the density of assignment variable.

In order to test if the first implication is satisfied in this context, I exploit the information on age, gender, high-school grades, the region of origin to have a first graphical inspection on the results of such tests[16]. Additional tests on other covariates are reported in the formal analysis in the next subsection.

Figure D.10 show how the covariates move with respect to the ISEE. The graphs do not display sharp discontinuities around the thresholds. The right-hand side part of each plot displays a more noisy pattern since the number of observations tends to decrease as we move on the top of the social ladder. However, even if the variance increases, these variables seem to display no change in mean around any of the thresholds.

Finally, to test whether any discontinuity exists in the density of assignment variable around each threshold, the McCrary test is performed [see McCrary (2008)]. This condition could fail if students could manipulate their ISEE in order to get a higher benefit. First, one must consider

[15]This evidence may be explained by a pure income effect. Since students on the right tail of the income distribution may have better-living conditions or more educated parents, than those at the bottom. Such reasoning is further strengthened by the fact that the benefit awarding is not conditioned on a minimum GPA, therefore, no direct effects of the benefits are expected.

[16]Notice that such information allows to observes students' relevant characteristics, which are likely to affect current academic performance.

that, in this context, the ISEE indicator must be certified by a professional agency and that violations are legally punished. The data confirm that, in fact, such possibility is not undertaken by the individuals in the sample. Figure D.11 summarizes the results. As it possible to notice, observations evolve smoothly around thresholds, so that there is no a systematic strategic manipulation of the income declaration. The only exception is at the third threshold. In particular, at a first look, one could think that the necessary condition for the RDD to hold might be violated, but as is shown in figure such result is mainly explained by differences in the composition of the sample across years. A possible additional explanation on such discontinuity may come from the fact that the frequency of families of income around 20.000 €sharply decreases at the population level, see figure D.13. Figure D.14 is the frequency distribution by classes of ISEE in the sample. As it possible to notice the sample and the population distribution are quite comparable.

3.4.2 Estimation Results

The regression results are reported in table D.3 where the estimates are computed for each threshold separately. The table shows the treatment effects of receiving different levels of benefit on credits. It is important to notice that all the benefits need not be reimbursed if some credit requirements are satisfied. Students aiming at fulfilling these requirements may study as for obtaining more credits, with potentially negative consequences on grades. Table D.3 reports results on GPA to check such hypothesis. These estimates confirm what the graphical inspection has firstly enlightened. Performance at the first two thresholds does not seem to change in response to a change of the benefit of around 800 €. Notice also that the mean number of credits in these samples are concentrated near the requirements.

On the other hand, the results on the linear prediction of performance around the third threshold are much stronger and significant. Those students who receive around 3500€more, perform actually worse than the students on the right-hand side of such threshold. In particular, they earned 9 credits less[17], around a 30% percent less, keeping a similar GPA. Pure credit requirements seem to be sufficient to increase students' performance in terms of credit achievement with no cost in term of GPA.

Taking this evidence together, we can say that those students, who are actually paying more than their counterparts, are incentivized to study more since they do not want to suffer an extra-cost given by delayed graduation, something which is in line with previous research [see Garibaldi et al. (2012)]. Notice that the results are stable even after including a set of controls into the regressions. This indicates that the potential problem of spurious regressions, which could arise from a global discontinuity in the data, should not be of much concern. Regarding the subsequent thresholds, it seems that no clear effects are in place. However, notice that as

[17] Which is comparable to one and a half course less.

we focus on wealthier students, namely those around the last thresholds, we get more noisy estimates and we see that the performance is centred at a lower level. The relaxation of the budget constraints could rationalize such results. Even if students at the right of the threshold may eventually pay an extra 10% more than those on the left, this does not induce them to put extra effort since they have enough own financing, i.e the budget constraint may be non-binding or slack.

Finally, notice that the local average treatment effects are estimated allowing for different linear polynomials on the two sides of each threshold. In the Appendix, D.4 evaluates the sensitivity of the above estimates with respect to the functional form of the control function in ISEE, $F(Y_i - c)$, included in equation 3.1. Specifically, second-degree polynomials are included and this robustness exercises support the validity of the main results reported in D.3.

Moreover, to further check the robustness of the results non-parametric estimates are reported in table D.5 of the Appendix, following the procedure in Calonico et al. (2014b) and in Calonico et al. (2014a).

On the whole, this quasi-experimental evidence confirms that students who are mostly entirely financed for their education take longer to complete their degrees, while those who are actually paying something for their education are relatively faster with no difference in terms of accuracy.

Sub-groups analysis

Regarding academic performance at higher education, it becomes quite important to distinguish between groups of subjects which have been shown to reach quite different achievements on this dimension [see for a review Buchmann et al. (2008)]. In order to gain some insights on the performance differentials in education, I investigate how the results vary with the gender dimension. The results are displayed in table D.6 and D.7 of the Appendix. Academic performance seems to differ substantially between female and male students, both in terms of credits and GPA. In particular, female students obtain a significantly higher amount of credits and better grades at almost all the thresholds of ISEE. Such results suggest that gender differentials in education are independent from the level of income of the family. Female students perform better than the males, without regards to their position in the social ladder. For what regards the sensitivity to incentives there is not much difference across genders at almost all the levels of initial wealth.

Another important dimension of heterogeneity, which is not captured by the average effects reported above, is given by the distance from graduation. In particular, given that students at different level of degree may suffer differently the costs of delayed graduation, it is important to distinguish between the graduate and undergraduate students. In particular, the former group has to reach indeed a lower total amount of credits in order to graduate with respect to undergraduate students[18]. This might impact the speed at which they proceed during each year. In particular, by having an extra year with respect to the graduate students, it could be that undergraduates have more time to smooth for catching-up with the lacking credits and potentially they suffer less the pressure of an extra-cost of delay at the first year of their studies. On the other side, it could be that undergraduates accumulate more unattained credits, since they have 60 credits more to obtain with respect to graduate students, and they eventually need more time to reach the amount needed for graduation. Which of these effects is anticipated by students is an empirical question. Table D.8 and D.9 of the Appendix show how bachelor students behave with respect to master's in terms of credits and GPA, respectively. Even across these different type of students, there is no a differential reaction to incentives at all the level of income. While as it is possible to notice from the amount of credits obtained, and most of all, from the GPA results, undergraduate students perform generally worse than graduates. Such results could be rationalized, as anticipated before, by the fact that they can smooth the credit load on a longer time interval, reducing, consequently the yearly effort, both in terms of quantity and quality. Other arguments could eventually provide a valid explanation for such results[19]. However, I don't have enough information to properly discriminate across explanations. Future research should address more deeply this issue.

3.5 Discussion

It is interesting to compare the overall findings with the results found in the literature. In particular, the evidence from Angrist et al. (2009) is in line with the findings from this study. The authors show that in a randomized experiment done with low-performer students[20], there are no incremental effects on GPA of pure financial incentives. Moreover, also in the Canadian context gender differentials were found. The results from this study look at the incentive effects of all type of students (low and high ability), for a wide range of the income distribution and in different years of observation. Even in such a wider sample, the results on the scholarship recipients confirm that if students do not suffer any cost for their education, it is unlikely that, by giving a higher benefit, any performance change is detected.

The results are also in line with Belot et al. (2007). In that study, it has been estimated the effect of reducing the maximum duration of a Dutch grant and the authors found positive results on performance. Such evidence is in line with the idea that studying effort react positively to an increase in the extra-cost of a delayed graduation.

In the study by Leuven et al. (2010), high-ability students who receive a high or low reward from passing all the exams in the first year of education, perform better than low-ability

[18]Undergraduates need to obtain 180 credits while graduate students 120.

[19]For example, it could be that the grading system followed by professor changes across degree levels

[20]The authors excluded those with a high-school GPA in the upper quantile

students who receive a high or low reward. Moreover, the low-ability students who receive the highest benefit perform significantly less than the same counterpart in the control group. Given that the scholarship, in this case, is conditioned on passing all the exams, it might be that the results are driven by the hardness of the requirement *per se* and not by a different reaction to monetary incentives. In the present study, given that requirements may be not binding, the problem of having this confounding mechanism is reduced.

The effects presented in the previous section add also new insights for what concerns the empirical literature on the Italian context. The study which proposes a similar explanation to the one suggested in this paper is the one by Garibaldi et al. (2012). The authors, indeed, have shown that an increase in the tuition fee significantly reduces the probability of late graduation. Since they focused their attention on the performance of students enrolled in the last year of the degree, we are not completely sure that a pure cost argument rationalizes the evidence, since we do not know what happened during their career in previous years.

In Schizzerotto et al. (2012), the authors show that monetary incentives (topped-up to the "Right to Study" incentives) do not generate any performance differences for students in the province of Trento. The Grant 5B incentive analyzed in the above-cited study requires the students to obtain 50 credits out of 60 in order not to reimburse the benefit. Indeed the students obtaining the Grant 5B on average obtained more credits than the pool of students in this sample at the first year of their study, even if the incentives *per se* is ineffective. Such results suggest that a possible way to increment the credit performance of the scholarship recipients in this sample would be to simply demand more credits as minimum requirement. In the randomized experiment described by De Paola et al. (2012) run in Calabria with high-ability students, large and small rewards produce very similar effects on performance. However, the authors do not discuss how much such grants actually cover of the total cost of attending college. Not knowing the incidence of benefit makes the comparison with the present study difficult.

Finally, Mealli and Rampichini (2012) have shown that "Right to study" scholarships do not seem to be effective in generating different decisions of drop-out for poorest students. Such results further corroborate the presented analysis on the probability of having obtained zero credits, table D.10, which might be a good proxy for identifying dropped-out students.

3.6 Conclusions

This paper has studied the effects of different levels and types of monetary incentives on academic achievement.

The main conclusion from this paper is that the incidence of such incentives with respect to the budget faced by the individual is crucial in determining the success of the program. In particular, students who receive a lower benefit and, therefore, who are actually paying some share of the total cost of college attendance, surprisingly perform better than those whose cost is completely subsidized. Given that students face similar requirement for non-returning the benefits (independently from the level of the benefit received), the main explanation which rationalizes this fact is the direct effect of the benefit on the cost of attending college. A lower benefit, relatively to a higher aid, increases the cost of attending college, and, consequently, it motivates students to finish early in order not to pay the extra costs of a delayed graduation.

Finally, interesting comparisons among different groups of subjects have shown striking evidence. Gender differentials in academic achievement persist considering a wide range of the income spectrum of families. Incentives seem to be not effective in reducing such gaps. Moreover, in line with the proposed explanation of the results, undergraduate students acquire fewer credits than master students, since they have more time to smooth in catching-up with the requirements for graduation. The threat of a possible extra-cost faced by such students at the first year is less compelling.

The results are also discussed with respect to the previous literature. Given the coverage of the sample and the richness of types of benefits analyzed in this paper, most of the previous findings could be rationalized in the present setting.

To conclude, the results may be particularly relevant for policy makers when evaluating the social costs and benefits of such monetary incentives, especially in Italy, where students performance are still behind those of scholars in other developed countries.

Bibliography

Aalto, S., U. Ayesta, and R. Righter (2009). On the gittins index in the m/g/1 queue. *Queueing Systems 63*(1-4), 437–458.

Allen, D. (2015). *Getting things done: The art of stress-free productivity*. Penguin.

Anderson, C. M. (2001). *Behavioral models of strategies in multi-armed bandit problems*. Ph. D. thesis, California Institute of Technology.

Angrist, J., D. Lang, and P. Oreopoulos (2009). Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics 1*(1), 136–163.

Arieli, A., Y. Ben-Ami, and A. Rubinstein (2011). Tracking decision makers under uncertaint. *American Economic Journal: Microeconomics 3*(4), 68–76.

Ariely, D. and K. Wertenbroch (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological science 13*(3), 219–224.

Arlotto, A., S. E. Chick, and N. Gans (2013). Optimal hiring and retention policies for heterogeneous workers who learn. *Management Science 60*(1), 110–129.

Becker, G. S. (1975). Front matter, human capital: a theoretical and empirical analysis, with special reference to education. In *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education, Second Edition*, pp. 22–0. NBER.

Bellman, R. (1957). Dynamic programming.

Belot, M., E. Canton, and D. Webbink (2007). Does reducing student support affect scholastic performance? evidence from a dutch reform. *Empirical Economics 32*(2-3), 261–275.

Bettinger, E. (2007). How financial aid affects persistence. *College Choices: The Economics of Where to Go, When to Go, and How to Pay for It*, 207.

Bray, R. L., D. Coviello, A. Ichino, and N. Persico (2016). Multitasking, multiarmed bandits, and the italian judiciary. *Manufacturing & Service Operations Management 18*(4), 545–558.

Buchmann, C., T. A. DiPrete, and A. McDaniel (2008). Gender inequalities in education. *Annu. Rev. Sociol 34*, 319–337.

Buser, T. and N. Peter (2012). Multitasking. *Experimental Economics 15*(4), 641–655.

Calonico, S., M. Cattaneo, and R. Titiunik (2014a). Robust data-driven inference in the regression-discontinuity design. *Stata Journal 14*(4), 909–946.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014b). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica 82*(6), 2295–2326.

Covey, S. R. (1991). *The 7 habits of highly effective people*. Simon & Schuster New York, NY.

Coviello, D., A. Ichino, and N. Persico (2014). Time allocation and task juggling. *The American Economic Review 104*(2), 609–623.

Coviello, D., A. Ichino, and N. Persico (2015). The inefficiency of worker time use. *Journal of the European Economic Association 13*(5), 906–947.

Crawford, S. and V. C. Wiers (2001). From anecdotes to theory: A review of existing knowledge on human factors of planning and scheduling. *Human performance in planning and scheduling*, 15–43.

Crenshaw, D. (2008). *The Myth of Multitasking: How "Doing It All" Gets Nothing Done*. John Wiley & Sons.

De Paola, M. and F. Gioia (2016). Who performs better under time pressure? results from a field experiment. *Journal of Economic Psychology 53*, 37–53.

De Paola, M., V. Scoppa, and R. Nisticò (2012). Monetary incentives and student achievement in a depressed labor market: Results from a randomized experiment. *Journal of Human Capital 6*(1), 56–85.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature 47*(2), 315–372.

Dynarski, S. and J. Scott-Clayton (2013). Financial aid policy: Lessons from research. Technical report, National Bureau of Economic Research.

Dynarski, S. M. (1999). Does aid matter? measuring the effect of student aid on college attendance and completion. Technical report, National bureau of economic research.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, 25–42.

Fryer, R. G. et al. (2011). Financial incentives and student achievement: Evidence from randomized trials. *The Quarterly Journal of Economics 126*(4), 1755–1798.

Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *The American Economic Review*, 1043–1068.

Gans, N., G. Knox, and R. Croson (2007). Simple models of discrete choice and their performance in bandit experiments. *Manufacturing & Service Operations Management 9*(4), 383–408.

Garibaldi, P., F. Giavazzi, A. Ichino, and E. Rettore (2012). College cost and time to complete a degree: Evidence from tuition discontinuities. *Review of Economics and Statistics 94*(3), 699–711.

Gittins, J., K. Glazebrook, and R. Weber (2011). *Multi-armed bandit allocation indices*. John Wiley & Sons.

Gittins, J. C. and D. M. Jones (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika 66*(3), 561–565.

Gneezy, U., S. Meier, and P. Rey-Biel (2011). When and why incentives (don't) work to modify behavior. *The Journal of Economic Perspectives 25*(4), 191–209.

Hahn, J., P. Todd, and W. Van der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica 69*(1), 201–209.

Horowitz, A. D. (1973). *Experimental study of the two-armed bandit problem*. Ph. D. thesis.

Johnson, E. J., C. Camerer, S. Sen, and T. Rymon (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory 104*(1), 16–47.

Johnson, E. J., J. W. Payne, J. R. Bettman, and D. A. Schkade (1989). Monitoring information processing and decisions: The mouselab system. Technical report, DTIC Document.

Johnson, G. E. (1990). Work rules, featherbedding, and pareto-optimal union-management bargaining. *Journal of Labor Economics*, S237–S259.

Kaelbling, L. P., M. L. Littman, and A. W. Moore (1996). Reinforcement learning: A survey. *Journal of artificial intelligence research 4*, 237–285.

Leuven, E., H. Oosterbeek, and B. Klaauw (2010). The effect of financial rewards on students'achievement: Evidence from a randomized experiment. *Journal of the European Economic Association 8*(6), 1243–1265.

Maule, A. J., G. R. J. Hockey, and L. Bdzola (2000). Effects of time-pressure on decision-making under uncertainty: changes in affective state and information processing strategy. *Acta psychologica 104*(3), 283–301.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of econometrics 142*(2), 698–714.

Mealli, F. and C. Rampichini (2012). Evaluating the effects of university grants by using regression discontinuity designs. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 175*(3), 775–798.

Meyer, R. J. and Y. Shi (1995). Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management science 41*(5), 817–834.

Payne, J. W., J. R. Bettman, and E. J. Johnson (1993). *The adaptive decision maker*. Cambridge University Press.

Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Schizzerotto, A., L. Vergolini, and N. Zanini (2012). Merit-based financial aid to students from low-income families and its effects on academic performance.

Simon, H. and W. Chase (1988). Skill in chess. In *Computer chess compendium*, pp. 175–188. Springer.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, 99–118.

Sneyers, E., T. Agasisti, K. De Witte, and G. Graziosi (2016). The impact of scholarships on students performance: a study on five italian universities.

Steyvers, M., M. D. Lee, and E.-J. Wagenmakers (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology 53*(3), 168–179.

Thistlethwaite, D. L. and D. T. Campbell (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology 51*(6), 309.

Tice, D. M. and R. F. Baumeister (1997). Longitudinal study of procrastination, performance, stress, and health: The costs and benefits of dawdling. *Psychological science*, 454–458.

Wald, A. (1950). Statistical decision functions.

Wilson, R. C., A. Geana, J. M. White, E. A. Ludvig, and J. D. Cohen (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General 143*(6), 2074.

Appendices

Appendix A

Appendix A - Chapter 1

Table A.1: Summary Statistics

Variable	N	Mean	St. deviation	Min	Max
Score	87	14.52	6.21	3	29
Gender (*)	87	0.25	0.44	0	1
Right Answers	87	16.19	6.36	5	30
Wrong Answers	87	6.70	3.58	1	18
Eco. Courses	87	0.07	0.25	0	1
Risk	87	6.22	1.51	3	9
Impulsivity 1	87	0.24	0.43	0	1
Impulsivity 2	87	0.28	0.45	0	1
Impulsivity 3	87	0.25	0.44	0	1

Notes: *Score* indicate the sum of all points obtained in the two parts of the test. (*) Gender equal to zero indicate the female sex. Notice that gender imbalance in the subject pool reflects the class composition of the "Liceo Minghetti" and does not depend by sample selection. *Right Answers* indicates the average of points gained from correctly answer the tasks. *Wrong Answers* indicates the average of points lost from not correctly answer the tasks. *Eco.courses* indicates the fraction of subjects who have taken part in any economic course. *Risk* is a subjective measure of risk attitude, which answers the following question: "In general, are you a person ready to take risks, or do you avoid to take risks? Please indicate your answer on a scale of 1 to 10, where 1 means that you do not want to take risk and 10 means that you are ready to take risks". The three *impulsivity* measures are derived from the answers to the three tests in Frederick (2005).

Table A.2: Test on Random Allocation

	Baseline Mean	Fixed Sequence Mean	Fixed Time Mean	p-value(*)
Eco. Courses	0.03	0.10		0.2718
			0.06	0.5369
Risk	6.13	6.31		0.2824
			6.24	0.7097
Impulsivity1	0.33	0.23		0.4602
			0.21	0.4156
Impulsivity2	0.33	0.29		0.6978
			0.21	0.2789
Impulsivity3	0.33	0.23		0.4602
			0.25	0.6156
Score (part 1)	10	9.75		0.7077
			9.70	0.9879
Obs.	30	28	29	

Notes: (*) Two-sample Wilcoxon Mann-Whitney test. *Score (part 1)* measure the performance obtained by answering the first part of the "Admission Test".

Table A.3: Answering Measures

Variable	Mean	St. deviation	Min	Max
Lookups per question	2.528	1.081	1	9
Total Lookups	50.275	8.716	34	80
Time for missing	186.312	132.355	11.187	653.343
Time for wrong	168.289	63.487	72.984	413.094
Time for right	102.031	21.068	50.324	160.246
Time for missing logic	279.333	136.141	81.213	653.343
Time for missing verbal	77.774	34.340	26.211	127.389
Time for missing math	162.314	124.270	111.879	609.359
Time for wrong logic	167.030	103.737	60.437	600.531
Time for wrong verbal	79.456	36.540	30.812	198.586
Time for wrong math	197.183	84.269	57.795	469.688
Time for right logic	115.430	36.490	51.651	191.231
Time for right verbal	56.378	18.811	23.702	118.421
Time for right math	129.645	49.373	45.906	286.591
Average time	139.448	14.745	84.157	150.555
Average time logic	130.982	34.031	66.396	237.544
Average time verbal	133.106	30.733	53.721	240.067
Average time math	154.257	44.281	64.036	286.591

Notes: All the above information are related to student behavior in the first part of the test. *Lookups per question* is the number of times the student has looked at the same question. *Total Lookups* is the total number of times the student has looked to all the questions. *Switch answer* is the number of times the individual has switched the answer. *Time for missing* indicates the average time spent on questions for which the answer was not provided. *Time for missing logic, verbal, math* measure the same time as the previous variable for each section separately. *Time for wrong* indicates the average time spent on questions for which the final answer was wrong. *Time for wrong logic, verbal, math* measure the same time as the previous variable for each section separately. *Time for right* indicates the average time spent on questions for which the final answer was right. *Time for right logic, verbal, math* measure the same time as the previous variable for each section separately. *Average time* is the mean time spent per question. *Average time logic, verbal, math* measure the same time as the previous variable for each section separately. Notice that all the time measures are expressed in seconds.

Table A.4: Mean number of lookups by categorization

Tot_Lookups > 51	Mean Number Lookups
0	2.391
1	3.239

Table A.5: Lookups categories

Tot_Lookups > 51	Freq.	Percent
0	73	83.91
1	14	16.09
Total	87	100.00

Table A.6: Mean Lookups by groups and parts for the "Switchers"

	Part	
Group	1	2
Baseline	3.235	3.129
Fixed Sequence	3.241	1.794
Fixed Time	3.242	3.000

Table A.7: Good and Bad Time-Allocation

Badtime	Freq.	Percent	Cum.
1	65	74.71	74.71
0	22	25.29	100
Total	87	100	

Table A.8: Probability of correctly and wrongly answer the tasks across groups and parts

Group		Part First	Part Second
Unconstrained	Pr(Right)	0.65	0.65
	Pr(Wrong)	0.26	0.27
Fixed Sequence	Pr(Right)	0.64	0.67
	Pr(Wrong)	0.24	0.22
Fixed Time	Pr(Right)	0.63	0.49
	Pr(Wrong)	0.25	0.33

Table A.9: Distribution of students according to their impulsivity

Impulsivity (categorical)	Freq.	Percent	Cum.
3	44	50.57	50.57
2	23	26.44	77.01
1	16	18.39	95.4
0	4	4.6	100
Impulsivity (indicator)			
1	44	50.57	50.57
0	43	49.43	100
Total	87	100	

Notes: The categorical variable shows the share of impulsive students in all the three questions of the Frederick'test. The category equal to three indicates that the student is impulsive in all the questions while the category equal to zero indicates that the subject was never impulsive. Notice that the indicator variable groups together all the students who, at least, were not impulsive in one question.

Table A.10: Regression estimates based on Impulsivity classification

	(1)	(2)
	Score	Δ Score (part2 -part1)
$FixSequence$ x $Part_2$ x $Impulsive$	3.100***	
	(1.140)	
$FixTime$ x $Part_2$ x $Impulsive$	1.358	
	(1.444)	
$FixSequence$ x $Part_2$	-1.343	
	(1.141)	
$FixTime$ x $Part_2$	-3.093**	
	(1.288)	
$FixSequence$	-0.298	-1.973*
	(0.693)	(1.013)
$FixTime$	1.348	-4.569***
	(0.707)	(1.203)
$FixSequence$ x $Impulsive$		4.661***
		(1.573)
$FixTime$ x $Impulsive$		4.069**
		(1.615)
$Part_2$	1.914**	
	(0.820)	
$Impulsive$	-1.823***	-4.708***
	(1.144)	(1.098)
$Notimpulsive$	10.821***	2.458***
	(0.535)	(0.695)
N	87	87
adj. R-sq	0.311	0.287

Notes: *Score* is the outcome variable, which sums the points obtained from the correct answers minus the penalty encountered from the wrong answers. *Intercept* represent the score obtained by the non-impulsive subjects in the first part of the test. *Impulsive* is a dummy variable equal to one if the subject was identified as impulsive in all the three questions of the Frederick'test. *Fixed Time* and *Fixed Sequence* are the dummy variable indicating if the subject belong to the respective treatment group. $Part_2$ it is a dummy variable equal to one if the score in the second part of the test is considered. Standard errors in parentheses. Significance levels: $*p < 0.10, **p < 0.05, ***p < 0.01$.

Table A.11: Lookups and Score by type and impulsivity

Types	Impulsivity	
	1	0
"Rational"		
Lookups	42.12	39.59
Score part 1	9.458	10.480
"Switchers"		
Lookups	61.25	49.25
Score part 1	4.75	10.75
"Bad Time"		
Lookups	40.71	38.32
Score part 1	9.201	11.065
"Switchers & Bad Time"		
Lookups	48.5	49.2
Score part 1	8.041	9.901

Appendix B

Appendix B - Chapter 1

Instructions

Welcome!
Thanks for your participation. This study aims to understand how people make decisions. The study was funded by the University of Bologna.
As a show-up fee you earn 8€. During the session, you can earn more money depending on your choices. Your earnings will be expressed in points and they will be converted into euros at the rate of 2 points = 1 euro. Payment will be made at the end of the session and privately.

At the end of the session, please remain seated. The researcher will come to your workstation to hand you a private envelope containing your payment.

During the study we do not allow any communication with other participants. You should turn off your phone now. For any question, please raise your hand and we will come to your workstation.

In this study, you will face a training test, whose structure and evaluation resemble the ones of the admission test to the Faculty of Economics and Statistics.

Now we will start to describe in detail the structure of the test and how you can use the computer in front of you to answer it.

<u>How the test is structured:</u>

The test is divided in two parts: Part 1 and Part 2.
Each part is composed by three thematic areas: Mathematics, Logic, and Verbal Comprehension.

The table below illustrates in detail the content of the test.

Table 1: Structure of the Test

Test Part 1				
Section	Number of tasks	Time available	Minimum Score	Maximum Score
Logic	6		-1.5	6
Verbal Comprehension	5		-1.25	5
Math	6		-1.5	6
Total	17	45	-4.25	17

Test Part 2				
Section	Number of tasks	Time available	Minimum Score	Maximum Score
Logic	6		-1.5	6
Verbal Comprehension	5		-1.25	5
Math	6		-1.5	6
Total	17	45	-4.25	17

Each part includes 17 questions, divided into the three sections as shown in Table 1. In particular:

- The first section contains 6 questions of logic.
- The second section includes 5 questions of verbal comprehension. This section requires to read a text on a generic topic. The text is followed by a series of questions, whose answers must be inferred exclusively from the content of the text.
- The third section contains 6 questions of math.

The questions are structured as multiple choices tasks. Each question has 5 possible answers, of which only one provides the correct solution.

For each question, it is assigned 1 point for correct answer, 0 points for missing and -0.25 points for wrong.

Table 1 shows the minimum and maximum possible score for each Part and for each section.

We will move now to a practical example that will show you how the test will appear on your computer and how you can select / de-select questions and answers.

Practical example:

As you can see, the first screen shows the questions of the logic section.

In general, to understand which section or which question is “active” you just have to see which of the sections’ box is coloured and what number is enlightened among the array of buttons on the bottom left of the screen.
In the center of the screen, you will see the number of the task and the text of the question. There are 5 multiple choices for answering each task, and each option is associated with a round selector.
In this example, we assume that you think that the right answer for question 1 is the second option. To select it, you have to position within the associated selector and click with the left button of the mouse. As you can see once you have clicked into it, the round selector and the button of the task became coloured. This indicates which question you answered and which answer you have chosen. To change your answer, just click inside the selector associated with the new option you want to select. Try to select answer 5. To clear the answer, just click back into the selector of the previously selected answer. Try to clear answer 5. You can now choose a new answer or choose to do not answer to this question. Suppose you choose to omit the answer for question 1, then go to question 3. As you can see, the box of task 1 is uncoloured, precisely because you have chosen to omit the answer.

You can choose to go back and forth to check the answers to all the questions of the various sections.
Let us now turn to the Verbal Comprehension section. Click on the button of the section. As you can see, the section box is coloured now. You can see the text to read and a ”T” in the bottom left corner, indicating that you are in the text screen.
Also in this section, as well as in Mathematics, you can answer to questions as explained in the previous example for the Logic section.
Now look at the time bar. In the right part of the screen there is a time scrolling bar which will be always displayed during the test and which tells you the minutes and seconds remaining to the end of the test. The bar has a mark every ten seconds: 60 seconds, 50 seconds, 40 seconds and so on. This is to show approximately how many seconds you have left.
Write now, on the instruction sheet, the remaining time at this moment.

Time remaining: .... minutes and .... seconds

<u>Please note:</u>
At the last minute, it will appear that the remaining minutes are zero but notice that you will still have 60 seconds to use, which will scroll down as time passes.

In Figure 2, you can see that about 15 seconds are left.

Figure 2

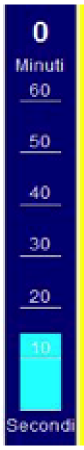

At any time you can click on the “End” button which will end the test and will bring you to the final results page. The test is also stopped automatically at the end of the available time.

Press now the “End” button.
We finished the practical example.
This study consists of two parts. Now we read the instructions relating to the first part. Please, pay attention.

Part 1

In this first part you have 45 minutes to answer 17 questions.
The questions are organized following the structure presented above. During the test you can move back and forth, both within and across sections and you can freely decide which questions you want to answer first and how much time you want to spend on each task, with a total available time equals to 45 minutes.
The final score is the sum of points earned by correctly answer the task minus the penalty for wrong answers, as shown above.

If there are any questions please raise your hand and we will come to your workstation.

Now we will start the first part of the study.

Part 2 - Fixed Time

The second part of the study is identical to the first one with just one difference.

In the first part you were free to decide in which sequence you want to answer the questions and how much time you want to spend on each of them.

In this part of the study you will only be able to freely decide in which order you want to answer the questions but you will have a maximum time for each question of 2 minutes and 22 seconds. In particular, you can still choose to switch freely among questions, even among the various sections, but you can not choose how much time you want to allocate to each question.

In particular, if you change question before the time for that task is finished, the timer will stop exactly when you switched the question; in this way you will not lose the remaining seconds and you can re-use them as soon as you return to the same question.

When the time for one question ends, you will see a message on the screen asking to choose another task.

The score is always computed in the same way: 1 point for each correct answer, -0.25 for each wrong answer, 0 for missing.

If there are any questions please raise your hand and we will come to your workstation.

Now we will start the second part of the study.

Part 2 - Fixed Sequence

The second part of the study is identical to the first one with just one difference.

In the first part you were free to decide in which sequence you want to answer the questions and how much time you want to spend on each of them.

In this part of the study you will be able to decide freely just how you want to allocate the time across questions while the sequence of questions is fixed. In particular, you can not choose to switch freely back and forth from one question to another, but you have to answer the test following a given order of the tasks: from task 1 to task 2 to task 3, starting from the Logic section, then moving to the Verbal and finally to the Math. Once you choose a section, you have to answer the tasks of that section in order of ascending numbers.

The score is always computed in the same way: 1 point for each correct answer, -0.25 for each wrong answer, 0 for missing.

If there are any questions please raise your hand and we will come to your workstation.

Now we will start the second part of the study.

Questionnaire

We now kindly ask you to complete this questionnaire.
The answers in this section will not affect your final score.
Some of these questions are related to personal information that will be useful for this study.
Your identity will not be revealed under any circumstances.
Please answer carefully.
Once you answered, you can not edit the answer.
Press OK to start.

Thank you!

1 Sex (Press the related button)

2 Age (Move your red triangle and press OK to confirm)

3 Have you ever attended economics courses?

4 Have you ever attended statistics courses?

5 Have you ever taken part in other researches within the university? (Select one of the answers and click OK)

6 In general, would you say that you can trust most people or that you are never too careful in dealing with people? (Select the answer and click OK to confirm)

7 In general, are you a person ready to take risks, or not? Please indicate your answer on a scale of 1 to 10, where 1 means ”I always prefer to do not take risks” and 10 means that risks ”I am always ready to take risks.”

8 In general, do you think it is important to help others, take care of their welfare? Please indicate your answer on a scale of 1 to 10, where 1 means ”not important at all” and 10 being ”surely important.”

Cognitive Reflection Test:

1 A bat and a ball cost $ 1.10 in total. The bat costs $ 1.00 blackberries than the ball. How much does the ball cost?

2 If it takes 5 machines 5 minutes to make five widgets, how long would it take 100 machines to make 100 widgets?

3 In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

Sample of questions from the Admission Test

B.1 Logic

You are the coach of the team that this year will represent the school at the Mathematics Olympics and you have to choose the last component of the team, between Mario, Franco, Michele, Jonathan and Vincenzo, selecting the one with better academic performance. Knowing that Mario has higher performance in school than Franco but not better than Michele. Jonathan has worse grades than Vincenzo, but not worse than Michele, who is selected?

A Michele

B Jonathan

C Mario

D Franco

E Vincenzo

B.2 Verbal Comprehension

B.2.1 Text

The Italian industry, at the moment of the unification, had diversified characters. The most important sector of northern Italy, the textile industry, was characterized by a widespread presence in rural areas of the territory. The close relationship between the agriculture sector and the industry is very evident in the case of the silk industry: the largest industry by number of nineteenth-century Italian workers. In this sector, the production of the raw material was not a mere sub-sector of agriculture. Rather it was the biggest source of wealth for the land owners and a means of poverty alleviation for many settlers. The second important textile industry, the cotton industry, differently from the silk sector, had very little productive organization of the weaving phase. The production was controlled by traders and entrepreneurs, who used domestic peasant labor. They were rare cases where the hand-loom work was concentrated in one room and few were the factories equipped with mechanical looms. Mechanization was instead become quite common since 1830/40 in the spinning sector. The wool industry was characterized by an underdeveloped structure both in spinning and in weaving. Excluding few mechanized factories, in the field dominated the home production. Again so you can see the close relationship between the agricultural class and the manufacturing work. This is even clearer in the case of the flax and of the hemp: apart from some mechanical flax spinning in Lombardy, the sector was not really industrialized. The steel industry was virtually non-existent, a part from a set of small forges (technologically underdeveloped) that were present in some areas of the territory. The situation was not much better for the mechanical industry. Apart from few large factories, the sector was made up of modest laboratories with no productive specialization, or simply handicraft; the same considerations apply to other areas, such as paper or the ceramics. In general, we can say that in northern Italy, there were concentrated almost all the Italian industries using mechanized production and the vast majority of manufacturing activity. In central Italy, only Tuscany was characterized by some noteworthy industries: the wool industry in Prato and the mining and metallurgical industries. Very few and underdeveloped were the industries located in southern Italy and in the islands. The few factories equipped with some machinery were mostly in Campania and precisely in Caserta, near Salerno and Naples. In the first two cities few cotton factories were founded by foreign entrepreneurs, while in Naples there were machine shops and metallurgical state-owned factories. From this summary we understand how, by the 1861, the Italian industry had already created its key development areas, namely the same as the major today production sectors. However, the Italy's industrial destiny was not yet marked decisively. What is certain is that the economic growth of the post-unification period gave priority to those areas where there were previous industrial structures and it neglected those parts of the country which did not experienced any substantial manufacturing production, or where such production was too small to stand the affirmation of the capitalist system.

B.2.2 Question

The subsequent development of the Italian industry has been characterized by:

A The manufacturing industry growth in the, previously, unaffected areas

B The manufacturing growth determined by the spread of the capitalist system

C The change in the pre-unification industry characteristics

D The consolidation of development gaps between different areas of the country

E A more balanced geographical distribution of the activities

B.3 Mathematics

The height of an equilateral triangle is equal to $\sqrt{3}$. The area A and the perimeter P of the triangle are equal respectively to:

A $A = \frac{1}{2}\sqrt{3}$; P=6

B $A = \sqrt{3}$; P=3

C $A = \sqrt{3}$; P=$\sqrt{6}$;

D $A = \sqrt{3}$; P=6;

E $A = 6$; P=$\sqrt{6}$;

Appendix C

Appendix - Chapter 2

Figure C.1: Hazard Rates by Task

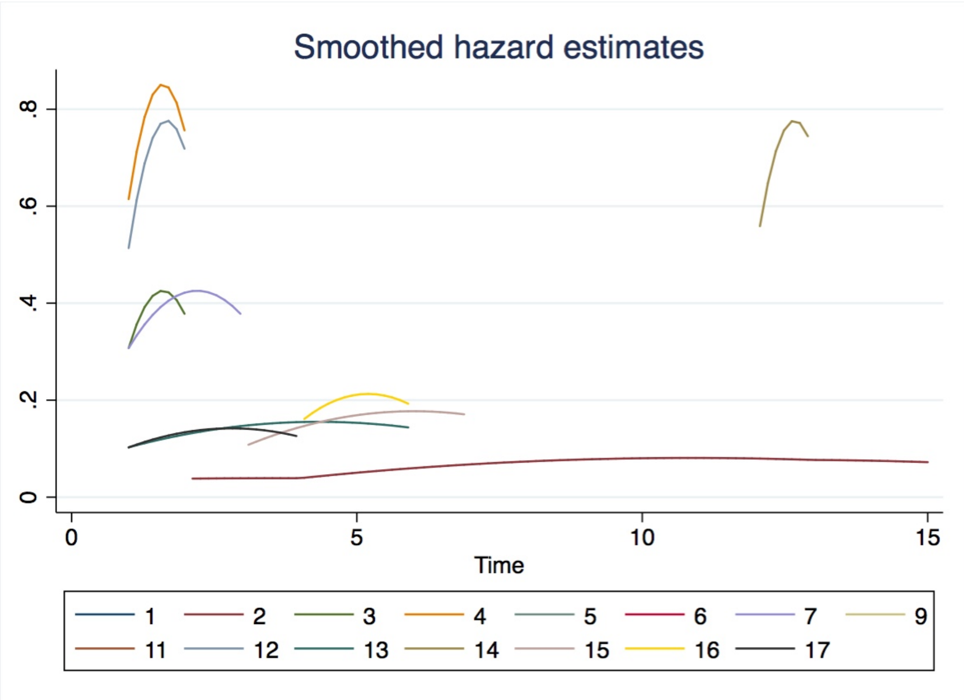

Appendix D

Appendix - Chapter 3

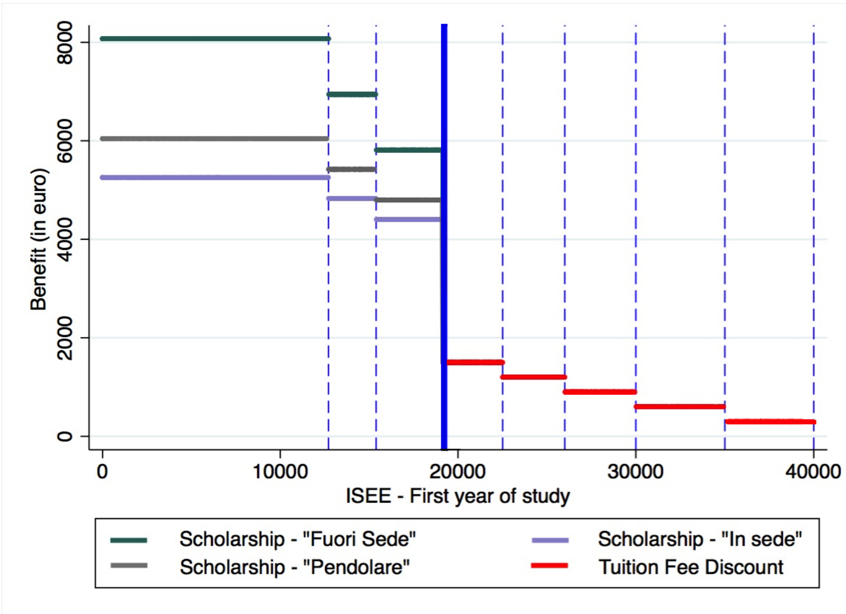

Figure D.1: Benefit Level Schedule - Scholarship and Tuition Fee Disocunt

Table D.1: Summary Statistics

Panel A	*Scholarship Recipients*							
Variable	Full Sample		First Threshold		Second Threshold		Third Threshold	
	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev
Percentage of Undergraduate students	0.672	0.469	0.688	0.463	0.662	0.473	0.647	0.477
Age	21,259	3,620	21,277	3,617	21,288	3,655	21,167	3,489
Undergraduate Grade	96,535	10,336	96,789	10,106	96,952	10,240	97,199	11,576
Area (frequency)								
Centre	21.92		20.91		20.06		20.89	
North	50.27		46.68		49.55		50.29	
South	20.38		24.18		22.17		20.77	
Islands	7.42		8.23		8.22		8.04	
High-School Grade	80,448	12,661	79,568	12,474	80,932	12,626	81,265	12,912
ISEE	15.121	2.213	12.685	773	15.331	760	18.122	1.900
ISPE	7.803	8.568	6.321	7.569	7.937	8.720	11.087	11.087
Credits	27,082	1,746	26,213	1,776	27,308	1,808	30,029	1,704
GPA	26,239	2,681	26,037	2,756	26,335	2,704	26,349	2,594
Gender	0.586	0.492	0.586	0.492	0.604	0.489	0.583	0.493
Observations	6103		1009		997		1704	

Panel B	*Tuition Fee Discount Recipients*							
Variable	Fourth Threshold		Fifth Threshold		Sixth Threshold		Seventh Threshold	
	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev	Mean	St.Dev
Percentage of Undergraduate students	0.590	0.494	0.653	0.478	0.765	0.427	0.352	0.492
Age	20,842	2,344	20,488	2,137	19,992	1,933	22,352	3,121
Undergraduate Grade	102,869	6,305	99,823	6,550	99.2	9,148	103,667	5,686
Area (frequency)								
Centre	21.85		25.26		25.10		25.28	
North	54.04		53.16		54.25		55.76	
South	17.70		16.18		15.79		14.50	
Islands	6.41		5.39		4.86		4.46	
High-School Grade	86,033	12,824	84,545	13,196	83,274	13,508	73,222	15,393
ISEE	22,810	2,039	25,511	2,103	30,221	2,267	32,814	2,452
ISPE	38,237	3,293	43,465	3,223	48,249	3,190	57,023	1,876
Credits	39,219	1,577	37,244	1,518	36,464	1,434	38,588	1,408
GPA	26,709	3,048	26,448	2,685	26,551	2,912	27,536	2,455
Gender	0.5	0.502	0.6	0.492	0.607	0.492	0.294	0.469
Observations	870		760		494		269	

Note: Statistics for the freshman who enrolled at University of Bologna in the 2009/2010 and 2010/2011 academic years by type of benefit and by threshold. Notice that the third threshold of statistics of Panel A group together the scholarship recipients and students who receive the 50% tuition fee discount.

Table D.2: Assignment of Tuition Fees Discount - 2010

ISEE Thresholds	Fess Discount
From €19152.98 to €22838	50%
From €22839 to €26390	40%
From €26391 to €30450	30%
From €30451 to €35525	20%
From €35526 to €40600	10%

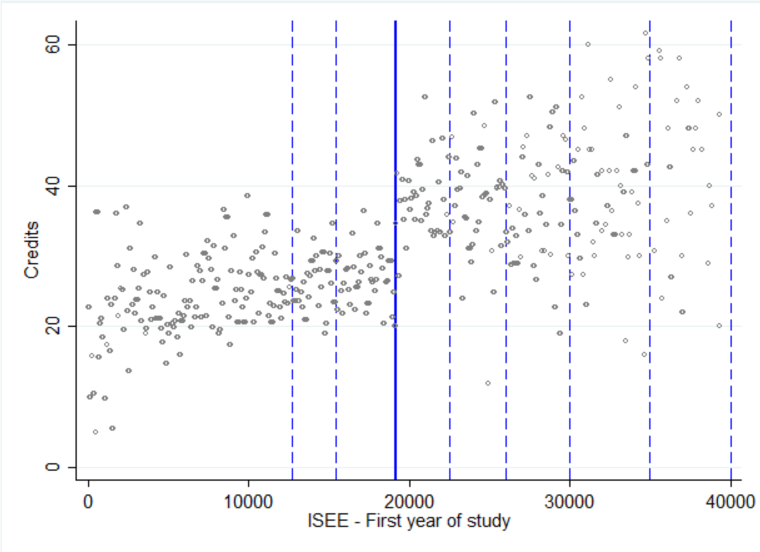

Figure D.2: Credits - 2009

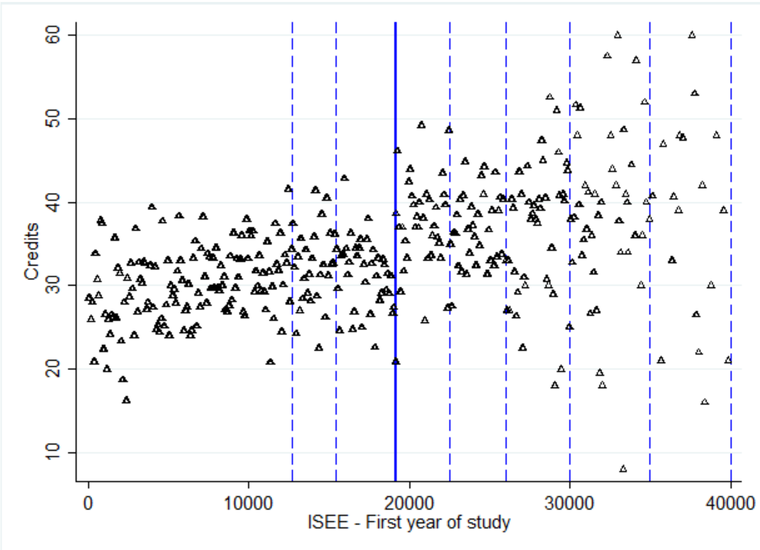

Figure D.3: Credits - 2010

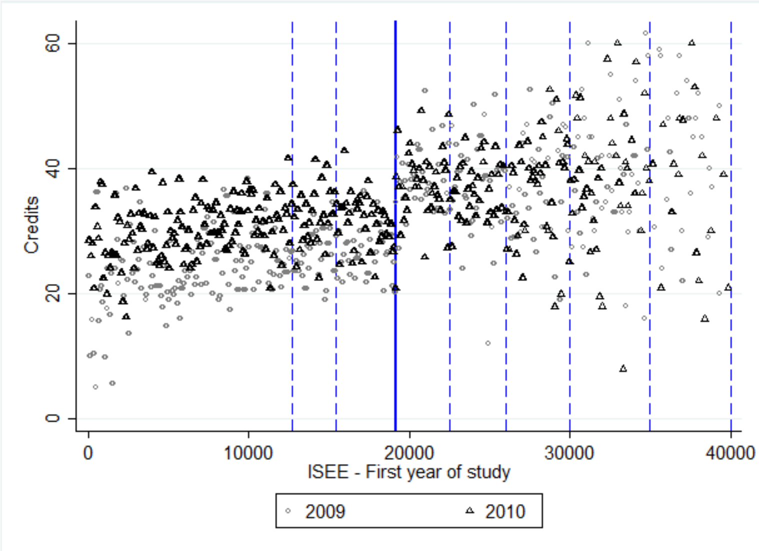

Figure D.4: Credits - 2009 and 2010

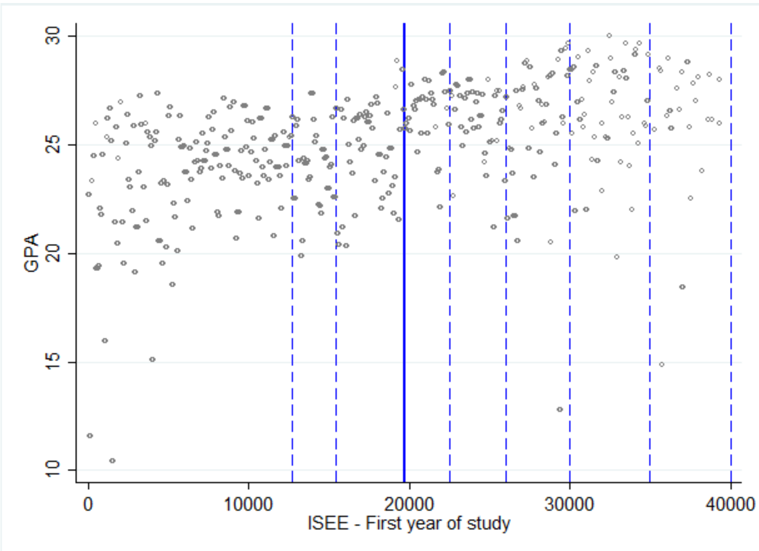

Figure D.5: GPA - 2009

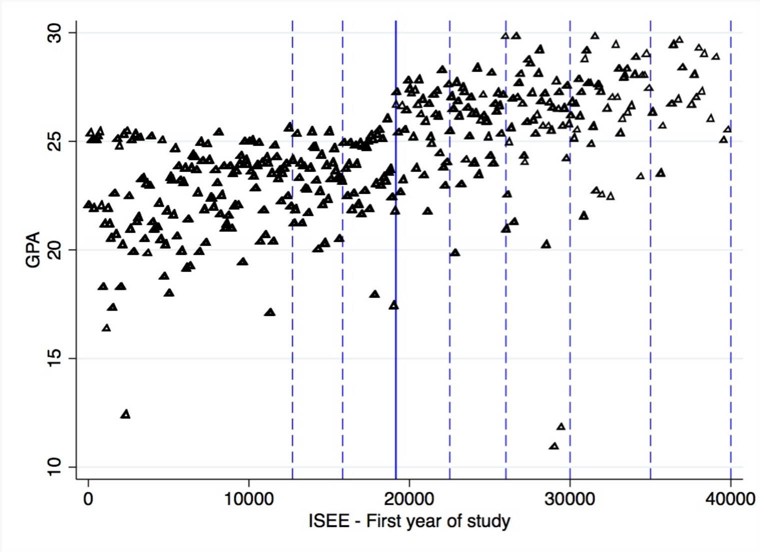

Figure D.6: GPA - 2010

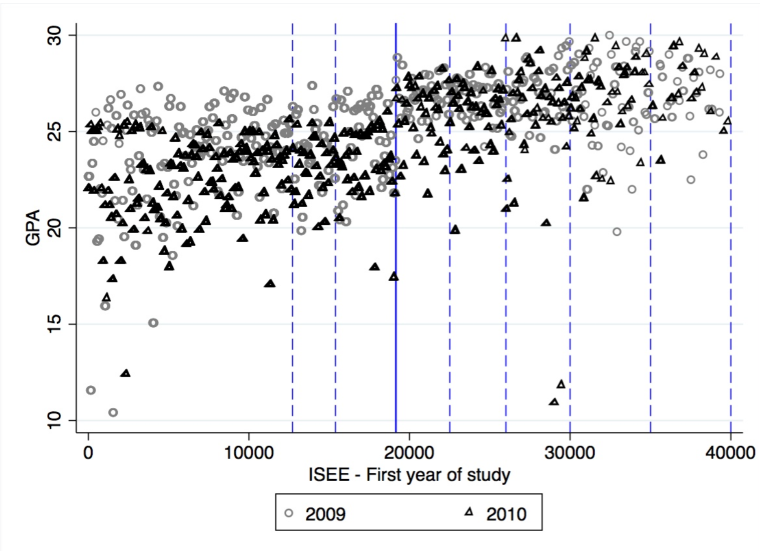

Figure D.7: GPA - 2009 and 2010

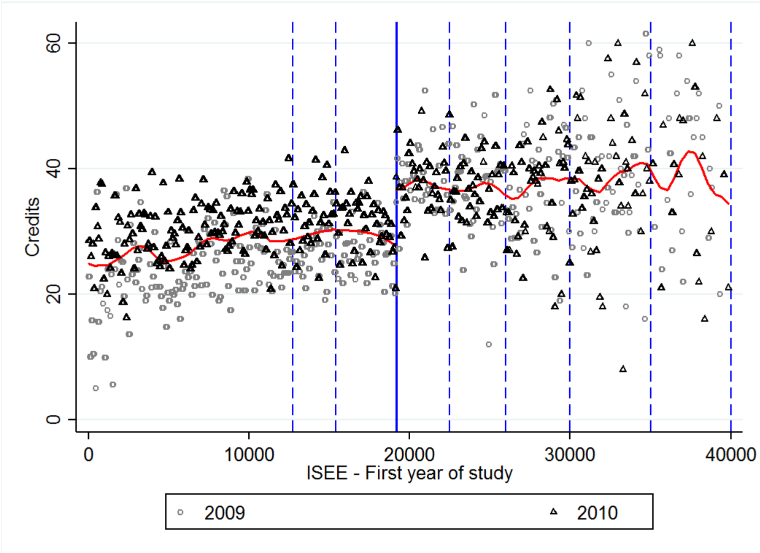

Figure D.8: Credits and Local Linear Regression - 2009 and 2010

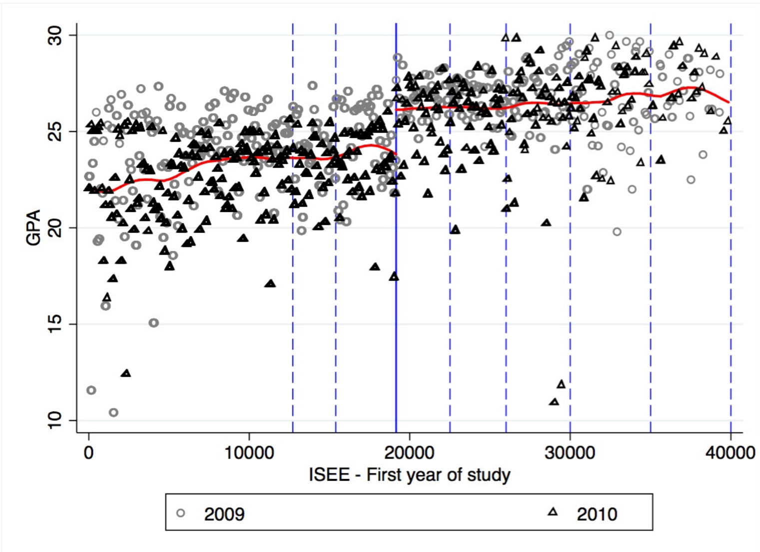

Figure D.9: GPA and Local Linear Regression - 2009 and 2010

Table D.3: Effects of benefits on Credits and GPA per threshold

	Credits		GPA	
Threshold 1: +800€	-1,347	-1,143	1,096	1,187
	(2.250)	(2.243)	(0.866)	(0.863)
Constant	30.552***	36.101***	23.753***	23.773***
	(1602)	(4.031)	(0.617)	(1.550)
Obs.	1013	1013	1013	1013
Threshold 2: +800€	1,009	1,112	0.753	0.779
	(2.278)	(2.272)	(0.886)	(0.882)
Constant	29.997***	28.715***	23.632***	23.565***
	(1.646)	(3.981)	(0.640)	(1.546)
Obs.	1003	1003	1003	1003
Threshold 3: +3500€	-8.460***	-8.402***	-1,158	-1,003
	(2.003)	(2.004)	(0.756)	(0.754)
Constant	36.605***	37.926***	25.761***	24.601***
	(1.760)	(3.321)	(0.664)	(1.249)
Obs.	1671	1671	1671	1671
Threshold 4: +300€	1,556	0.826	-0.369	-0.412
	(2.148)	(2.141)	(0.744)	(0.737)
Constant	35.941***	41.560***	26.136***	23.309***
	(1.463)	(3.819)	(0.507)	-1,315
Obs.	842	842	842	842
Threshold 5: +300€	4.132*	3,662	0.310	0.246
	(2.305)	(2.282)	(0.843)	(0.827)
Constant	33.734***	37.682***	25.240***	21.180***
	(1.697)	(4.008)	(0.621)	-1,453
Obs.	758	758	758	758
Threshold 6: +300€	4.328*	4.378*	0.068	0.197
	(2.542)	(2.517)	(0.926)	(0.885)
Constant	35.759***	28.651***	26.030***	17.828***
	(1.836)	(4.816)	(0.669)	(1.694)
Obs.	494	494	494	494
Threshold 7: +300€	3,817	3,626	2.142*	1,876
	(4.076)	(4.069)	(1.205)	(1.152)
Constant	38.270***	32.542***	25.172***	19.201***
	(3.352)	(6.499)	(0.991)	(1.839)
Obs.	268	268	268	268
Pre-treat. controls		Yes		Yes

Note. OLS estimates of linear equation 3.1 where Z_i is the vector of outcomes variables (Credits and GPA); $F(Y_i - c)$ is a vector whose elements are two polynomials (one for each side of the threshold) in the absolute difference between the ISEE declaration and the thresholds ; D_i is a dummy taking value 1 for students on the left side of the threshold. Robust standard errors in parentheses. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

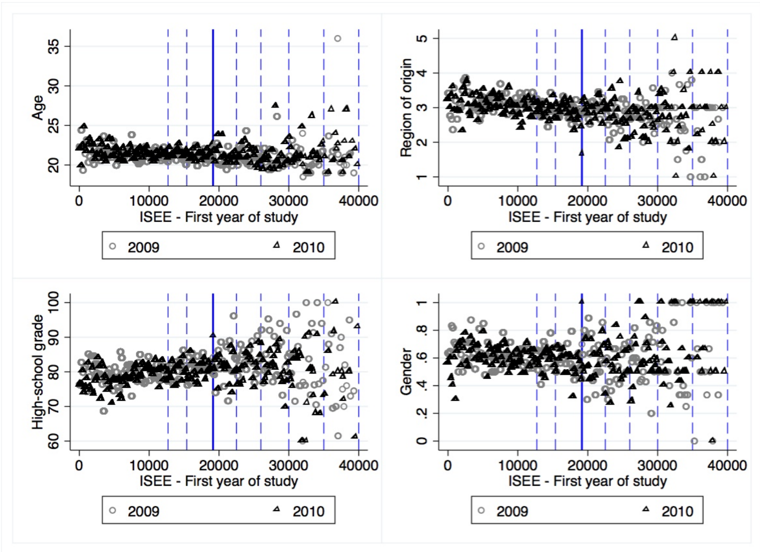

Figure D.10: Student characteristics - 2009 and 2010

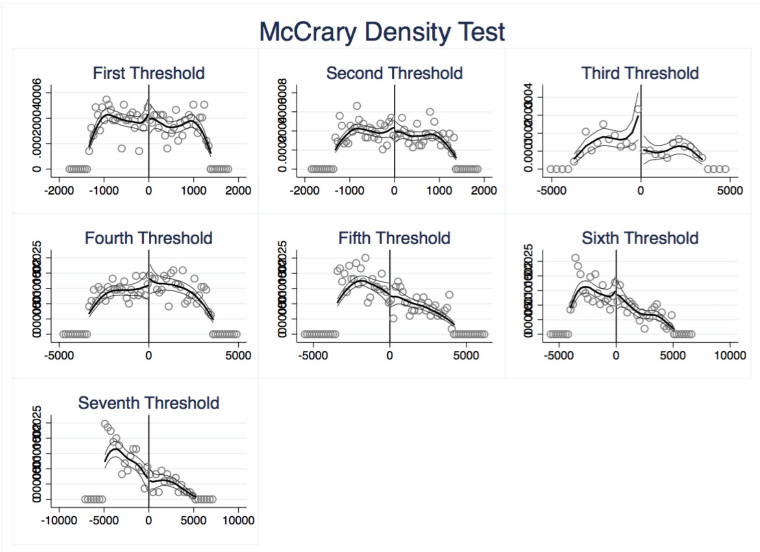

Figure D.11: McCrary density tests: individual thresholds

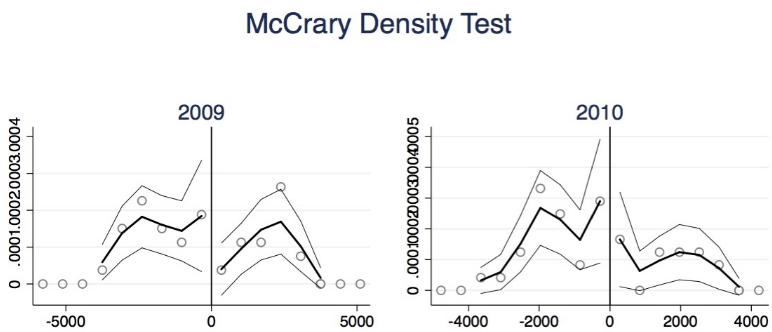

Figure D.12: McCrary density tests by year - Third threshold

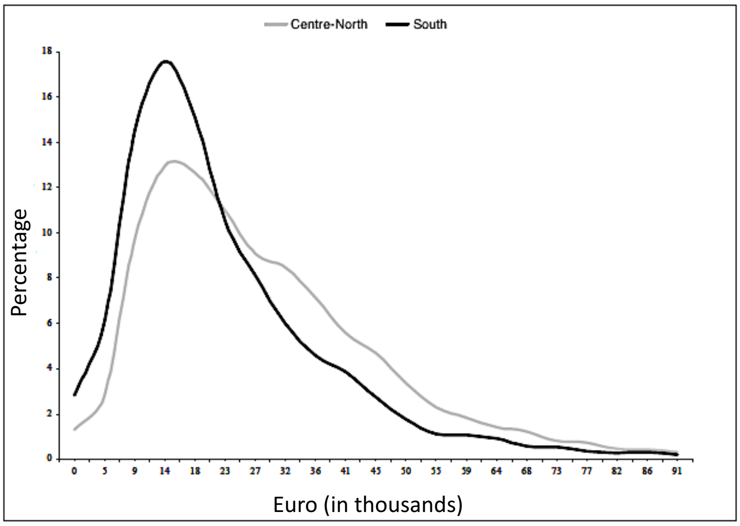

Figure D.13: Income distribution of the italian families
Elaboration from the Italian national institute of statistics report (EU-SILC) - ISTAT[1]

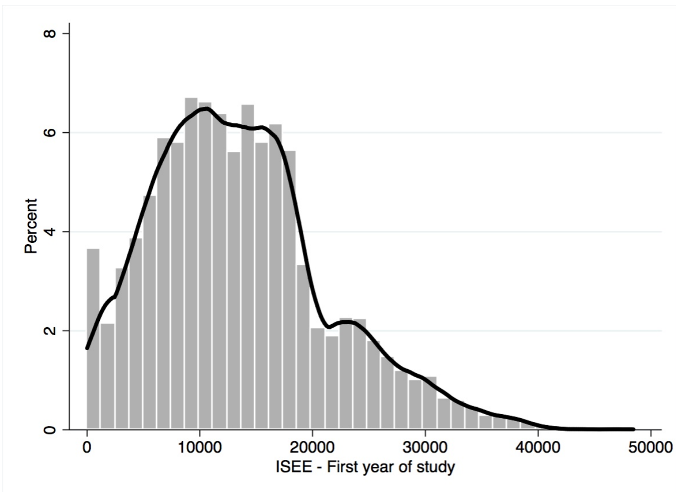

Figure D.14: Income distribution in the sample
Frequency distribution by bins of ISEE and estimated Epanechnikov kernel density.[Bin=39, start=0, width=1242.1026][2]

Table D.4: Effects of benefits on Credits and GPA per threshold - 2nd degree polynomials

	Credits		GPA	
Threshold 1: +800	-7.040**	-6.389*	-0.906	-0.544
	(3.278)	(3.276)	(1.262)	(1.261)
Constant	34.831***	40.273***	25.265***	25.193***
	(2.343)	(4.379)	(0.902)	(1.685)
Obs.	1013	1013	1013	1013
Threshold 2: +800	4,821	4,497	2.289*	2.159*
	(3.333)	(3.319)	(1.295)	(1.287)
Constant	28.917***	28.355***	23.617***	23.848***
	(2.416)	(4.339)	(0.429)	(1.683)
Obs.	1003	1003	1003	1003
Threshold 3: +3500	-9.657***	-9.465***	-2.188*	-1.972*
	(3.091)	(3.090)	(1.162)	(1.158)
Constant	35.891***	37.225***	25.184***	24.085***
	(2.727)	(3.943)	(1.025)	(1.478)
Obs.	1671	1671	1671	1671
Threshold 4: +300	1,475	0.656	-0.809	-1,147
	(3.168)	(3.152)	(1.097)	(1.085)
Constant	35.314***	41.059***	26.179***	23.498***
	(2.151)	(4.050)	(0.745)	(1.394)
Obs.	842	842	842	842
Threshold 5: +300	5,163	4,896	1,234	1,323
	(3.537)	(3.498)	(1.293)	(1.268)
Constant	31.635***	35.782***	24.374***	20.425***
	(2.569)	(4.384)	(0.939)	(1.589)
Obs.	758	758	758	758
Threshold 6: +300	-1,179	-1,055	-1,100	-0.736
	(3.681)	(3.654)	(1.344)	(1.289)
Constant	38.591***	32.233***	26.173***	18.090***
	(2.596)	(5.224)	(0.948)	(1.842)
Obs.	494	494	494	494
Threshold 7: +300	15.043**	15.501**	3.382*	3.534**
	(6.162)	(6.150)	(1.838)	(1.757)
Constant	31.256***	25.643***	24.256***	18.016***
	(4.855)	(7.295)	(1.448)	(2.084)
Obs.	268	268	268	268
Pre-treat. controls		Yes		Yes

Note. OLS estimates of 2nd degree polynomial equation 3.1 where Z_i is the vector of outcomes variables (Credits and GPA); $F(Y_i - c)$ is a vector whose elements are 2nd degree polynomials (one for each side of the threshold) in the absolute difference between the ISEE declaration and the thresholds ; D_i is a dummy taking value 1 for students on the left side of the threshold. Robust standard errors in parentheses. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

Table D.5: Effects of benefits on Credits and GPA per threshold - Non parametric

	(1)		(2)		(3)		(4)		(5)		(6)		(7)	
	Credits	GPA	Credits	GPA	Credits	GPA	Credits	GPA	Credits	GPA	Credits	GPA	Credits	GPA
Effect of higher benefit (conventional)	-12.63*	-3.44*	-1.91	0.26	-13.63***	-1.83	5.90	-0.343	6.87	2.63	6.94	2.12	10,421	2.68
	(6.73)	(1.83)	(4.58)	(1.55)	(4.12)	(1.25)	(4.13)	(1.18)	(4.43)	(2.28)	(4.83)	(1.25)	-(7.98)	(1.56)
Effect of higher benefit (bias_corrected)	-14.97**	- 4.09**	-3.08	-0.06	-15.399***	-1.89	7.42*	-0.26	7.58*	2.74	8.78*	2.65	9.67	2.42
	(6.73)	(1.83)	(4.58)	(1.55)	(4.12)	(1.25)	(4.13)	(1.18)	(4.43)	(2.28)	(4.83)	(1.25)	(7.98)	(1.56)
Effect of higher benefit (robust)	-14.97*	- 4.09**	-3.08	-0.06	-15.399***	-1.89	7.42	-0.26	7.58	2.74	8.78	2.65	9.67	2.42
	(7.76)	(2.02)	(5.43)	(1.85)	(4.74)	(1.52)	(4.71)	(1.33)	(5.43)	(2.89)	(5.60)	(1.55)	(9.64)	(1.58)
Bandwith for Loc. Poly (h)	268.32	262.72	352.39	420.01	798.47	1561.04	845.51	1382.28	1289.33	310.17	372.93	832.59	1424.77	2066.08
Bandwith for bias (b)	461.37	515,612	575.83	693.48	1314.56	2197.86	1430.39	2075.91	1948.37	751.46	695.67	1672.79	2196.89	3140.60
Number of observations	216	211	265	316	352	692	231	360	264	136	136	106	60	102

Note: Non-parametric estimates of the effect of higher benefit on performances: credits obtain and GPA at the end of the first year of college. I follow Calonico et al. (2014b) and I use their programmed command in Calonico et al. (2014a). Optimal bandwidth for the local polynomial (h) and for the bias (b) are reported. The treatment effect are reported for: the local polynomial estimator (conventional), the bias-corrected estimator proposed by Calonico et al. (2014b) and the same estimator with robust standard errors. The running variable is the distance of ISEE from the threshold value: $(Y-c)$. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

Table D.6: Effects of benefits on Credits by gender and thresholds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
females	3.644**	2,338	3.007***	3.170*	2.967*	4.560**	4.487**	4,309
	(1.548)	(1.643)	(1.127)	(1.707)	(1.534)	(1.864)	(2.147)	(3.578)
female * effect of higher benefit	-6.22*	2.10	-1.98	-1.30	-0.59	-0.379	-1.27	-6.01
	(3.38)	(3.51)	(2.43)	(2.53)	(3.61)	(3.66)	(4.18)	(5.86)
effect of higher benefit - Threshold 1	3,432							
	(2.870)							
effect of higher benefit - Threshold 2		-0.311						
		(3.034)						
effect of higher benefit - Aggregate Thresholds Scholarship			1,582					
			(2.086)					
effect of higher benefit - Threshold 3				-7.507***				
				(2.435)				
effect of higher benefit - Threshold 4					1,444			
					(3.094)			
effect of higher benefit - Threshold 5						4,187		
						(3.132)		
effect of higher benefit - Threshold 6							5,075	
							(3.626)	
effect of higher benefit - Threshold 7								7,374
								(5.295)
Constant (males)	27.290***	28.685***	27.979***	34.713***	34.557***	31.072***	32.916***	35.713***
	(1.762)	(1.849)	(1.276)	(2.000)	(1.625)	(2.010)	(2.274)	(3.931)
R-squared	0.004	0.003	0.004	0.040	0.002	0.011	0.027	0.026
Obs.	1010	1001	2011	1670	842	758	494	268

Note. OLS estimates of linear equation 3.1 where Z_i is the outcome variable Credits; $F(Y_i - c)$ is a vector whose elements are two linear functions (one for each side of the threshold) of the absolute difference between the ISEE declaration and the thresholds ; D_i is a dummy taking value 1 for students on the left side of the threshold. Robust standard errors in parentheses. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

Table D.7: Effects of benefits on GPA by gender and thresholds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
females	1.547**	1.273*	1.412***	1.426**	1.357**	2.056***	0.811	2.193**
	(0.628)	(0.659)	(0.454)	(0.656)	(0.530)	(0.679)	(0.781)	(1.063)
female * effect of higher benefit	-0.036	0.850	0.394	-0.285	0.570	-1.440	1,871	-2.943*
	(1.37)	(1.41)	(0.982)	(0.975)	(1.24)	(1.335)	(1.523)	(1.741)
effect of higher benefit - Threshold 1	1,120							
	(1.166)							
effect of higher benefit - Threshold 2		0.192						
		(1.218)						
effect of higher benefit - Aggregate Thresholds Scholarship			0.667					
			(0.841)					
effect of higher benefit - Threshold 3				-0.922				
				(0.936)				
effect of higher benefit - Threshold 4					-0.966			
					(1.068)			
effect of higher benefit - Threshold 5						1,033		
						(1.141)		
effect of higher benefit - Threshold 6							-1,119	
							(1.320)	
effect of higher benefit - Threshold 7								3.882**
								(1.573)
Constant (males)	22.855***	22.902***	22.876***	24.910***	25.503***	24.039***	25.516***	23.870***
	(0.715)	(0.742)	(0.514)	(0.769)	(0.561)	(0.732)	(0.828)	(1168
R-squared	0.006	0.003	0.007	0.012	0.008	0.014	0.020	0.016
Obs.	1013	1003	2016	1671	842	758	494	268

Note. OLS estimates of linear equation 3.1 where Z_i is the outcome variable GPA; $F(Y_i - c)$ is a vector whose elements are two linear functions (one for each side of the threshold) of the absolute difference between the ISEE declaration and the thresholds ; D_i is a dummy taking value 1 for students on the left side of the threshold. Robust standard errors in parentheses. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

Table D.8: Effects of benefits on Credits by level of degree and thresholds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Undergraduates	-2,571	-4.724**	-3.626***	-1,289	-0.385	-4.543**	0.190	-6.165*
	(1.729)	(1.832)	(1.257)	(1.767)	(1.660)	(2.025)	(2.350)	(3.664)
undergraduates * effect of higher benefit	-4,973	0.597	-2,142	-3,623	-2,265	2,971	-3,751	10.440*
	(3.945)	(3.831)	(2.734)	(2.732)	(3.767)	(4.001)	(4.580)	(6.03)
effect of higher benefit - Threshold 1	3,190							
	(3.660)							
effect of higher benefit - Threshold 2		1,404						
		(3.419)						
effect of higher benefit - Aggregate Thresholds Scholarship			2,173					
			(2.490)					
effect of higher benefit - Threshold 3				-6.009**				
				(2.674)				
effect of higher benefit - Threshold 4					2,756			
					(3.213)			
effect of higher benefit - Threshold 5						3,277		
						(3.548)		
effect of higher benefit - Threshold 6							7.467*	
							(3.971)	
effect of higher benefit - Threshold 7								0.965
								(5.664)
Constant (males)	31.526***	32.947***	32.260***	37.893***	35.943***	35.971***	33.720***	40.191***
	(2.017)	(2.022)	(1.427)	(2.064)	(1.848)	(2.254)	(2.539)	(4.483)
R-squared	0.008	0.015	0.013	0.044	-0.005	0.006	0.025	0.024
Obs.	918	906	1824	1493	760	669	426	236

Note. OLS estimates of linear equation 3.1 where Z_i is the outcome variable Credits; $F(Y_i - c)$ is a vector whose elements are two linear functions (one for each side of the threshold) of the absolute difference between the ISEE declaration and the thresholds ; D_i is a dummy taking value 1 for students on the left side of the threshold. Robust standard errors in parentheses. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

Table D.9: Effects of benefits on GPA by level of degree and thresholds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Undergraduates	-2.400***	-2.010***	-2.210***	-3.547***	-2.133***	-4.399***	-2.839***	-3.856***
	(0.678)	(0.720)	(0.493)	(0.651)	(0.535)	(0.710)	(0.828)	(1.071)
undergraduates * effect of higher benefit	1,232	0.296	0.769	1,152	-0.844	2,057	-1,788	2,848
	(1.552)	(1.508)	(1.075)	(1.005)	(1.214)	(1.403)	(1.613)	(1.762)
effect of higher benefit - Threshold 1	0.253							
	(1.437)							
effect of higher benefit - Threshold 2		0.527						
		(1.346)						
effect of higher benefit - Aggregate Thresholds Scholarship			0.395					
			(0.977)					
effect of higher benefit - Threshold 3				-2.107**				
				(0.984)				
effect of higher benefit - Threshold 4					0.154			
					(1.036)			
effect of higher benefit - Threshold 5						-0.716		
						(1.244)		
effect of higher benefit - Threshold 6							0.530	
							(1.399)	
effect of higher benefit - Threshold 7								0.185
								(1.655)
Constant (males)	25.324***	25.069***	25.193***	28.373***	27.325***	28.007***	27.813***	27.659***
	(0.785)	(0.796)	(0.558)	(0.760)	(0.596)	(0.790)	(0.894)	(1.310)
R-squared	0.026	0.028	0.030	0.042	0.064	0.066	0.115	0.121
Obs.	921	908	1829	1494	760	669	426	236

Note. OLS estimates of linear equation 3.1 where Z_i is the outcome variable GPA; $F(Y_i - c)$ is a vector whose elements are two linear functions (one for each side of the threshold) of the absolute difference between the ISEE declaration and the thresholds ; D_i is a dummy taking value 1 for students on the left side of the threshold. Robust standard errors in parentheses. Significance levels: * at 10%; ** significant at 5%; *** significant at 1% or better.

Table D.10: Effect of higher benefit at Threshold 3 on Probability of having zero credits

Treatment effects on Pr(Credits=0)		
Non parametric		
	Conventional	-0.71
		(0.50)
	Bias_corrected	-0.08
		(0.05)
	Robust	-0.08
		(0.06)
Bandwith for Loc. Poly (h)		1,208
Bandwith for bias (b)		1,735
Number of observations		534
Linear		0.042
		(0.028)
Constant		0.000
		(0.044)
2nd Degree Polynomials		0.074*
		(0.042)
Constant		0.025
		(0.053)
Number of observations		1674