Contributions

Maria Bigoni, Margherita Fort, Mattia Nardotto
and Tommaso G. Reggiani*

Cooperation or Competition? A Field Experiment on Non-monetary Learning Incentives

DOI 10.1515/bejeap-2014-0109

Published online June 25, 2015

Abstract: We assess the effect of two antithetic non-monetary incentive schemes based on grading rules on students' effort, using experimental data. We randomly assigned students to a tournament scheme that fosters competition between paired up students, a cooperative scheme that promotes information sharing and collaboration between students and a baseline treatment in which students can neither compete nor cooperate. In line with theoretical predictions, we find that competition induces higher effort with respect to cooperation, whereas cooperation does not increase effort with respect to the baseline treatment. Nonetheless, we find a strong gender effect since this result holds only for men while women do not react to this type of non-monetary incentives.

Keywords: competition, cooperation, education, experimental economics, gender differences, incentives

JEL Classification: A22, C93, I20

1 Introduction

The high incidence of problems of school dropout, student disaffection with school and low academic achievement motivates policy makers' growing interest

*Corresponding author: **Tommaso G. Reggiani**, Department of Business Ethics, University of Cologne, Albertus-Magnus-Platz, Cologne, Germany; IZA, Bonn, Germany, E-mail: tommaso.reggiani@uni-koeln.de

Maria Bigoni, Department of Economics, University of Bologna, Piazza Scaravilli 2, Bologna, Italy, E-mail: maria.bigoni@unibo.it

Margherita Fort, Department of Economics, University of Bologna, Piazza Scaravilli 2, Bologna, Italy; IZA, Bonn, Germany; CESifo, Munich, Germany, E-mail: margherita.fort@unibo.it

Mattia Nardotto, Department of Economics, University of Cologne, Albertus-Magnus-Platz, Cologne, Germany, E-mail: mattia.nardotto@uni-koeln.de

in means of increasing students' effort at every level of education. Achieving this goal is especially difficult now when the pressure to cut public spending is rising and most governments are facing constraints in their ability to provide public services, including education. Thus, promoting students' effort by means of effective non-financial incentives may become particularly attractive.¹

In this paper we study the effect of one specific type of non-monetary incentive, namely the way university students' work is evaluated by lecturers. We study the effect of different grading rules on effort by assigning students to alternative incentive schemes: a tournament, a piece rate and a scheme that promotes cooperation. We reward university students with extra points for their course grade: these additional points may result in a student passing the exam rather than failing, which in turn can have consequences on how long it will take that student to successfully complete his/her studies. We exploit the between-subjects design of our field experiment, where each subject is exposed to only one incentive scheme, in order to test the predictions of a model that contemplates all the incentive schemes mentioned above. The model predicts that in the competitive environment the effort should be higher than in the benchmark (the piece rate) while the effort under the benchmark should be higher than under the cooperative scheme. Our data confirm the theoretical predictions in the full sample. Moreover, we show that an important difference emerges between genders: promoting competition appears to have a strong positive effect on the exerted effort only for males. In contrast, promoting cooperation does not significantly reduce effort in either gender compared to when students can neither compete nor cooperate.² This paper complements recent studies that assess how school achievement is affected by financial incentives on input (e.g., subsidizing the purchase of learning supports) or on output (e.g., rewarding students if they pass the exams or obtain high grades) and represents a first attempt to compare cooperative and competitive incentive schemes based on non-financial rewards for students' performance. Our findings are relevant for the literature because non-financial incentives represent a relatively cheap way of increasing students' effort.

The paper proceeds as follows. After a brief review of the related literature (Section 2) we describe and discuss in detail our experimental design (Section 3). In Section 4, we present a simple model and derive the theoretical predictions

¹ As discussed in Section 2, studies on financial incentives have proved to be successful in improving students' performance; however, the cost of inducing higher effort is not negligible (see Fryer 2011).

² This finding has been recently confirmed by a further field study on relative grading systems by Czibor et al. (2014).

which will serve as a reference for the analysis of the experimental data presented in Section 5. Section 6 concludes and presents possible further developments of this research.

2 Related Literature

Our study relates to the empirical literature that examines how to foster students' effort and school achievement through explicit incentive schemes.

Most papers explore the role of financial incentives: some papers focus on college students, whereas others focus on younger students. Among the latter category, Blimpo's (2014) study is the closest to our experiment. Analyzing data from a field experiment in Benin with a pool of 100 secondary schools, he studies whether individual incentives or different kinds of team incentives can lead to a higher students' performance. He finds that an individual-based incentive scheme with cut-off target is most effective for students at an intermediate performance level. By contrast, when students work in teams, a tournament scheme yields the most beneficial effects, inducing all the teams' members to work harder. Kremer et al. (2009) focus on the evaluation of a merit scholarship program dedicated only to female students in elementary schools in Kenya. They observe a substantial increase in exam scores: in particular, girls with low pre-test scores, who were unlikely to win a scholarship, reported positive and significant gains in terms of higher school performance. Fryer (2011) assesses the impact of monetary incentives both on input (e.g., subsidizing the purchase of learning supports) and on output (e.g., giving money based on grades, or conditional on passing the exam). Results show that incentives on "inputs" can raise achievement even among the poorest minority students in the lowest performing schools. Incentives focused on "output" prove to be much less effective.

De Paola et al. (2012) study the effectiveness of financial incentive schemes in enhancing students' performance exploiting data of a randomized experiment on college students in an Italian university. They find that financial rewards in a tournament contribute to increase high ability students' performance but they are much less effective when it comes to lower-ability students. Similarly, Leuven et al. (2010) find evidence that the effects of financial rewards on the probability that freshmen at Amsterdam University pass all first year requirements vary depending on students' ability.

Some studies on financial incentives in education have revealed relevant gender effects. Angrist and Lavy (2009) evaluate the effectiveness of financial rewards on the achievement of Israeli student by means of a randomized

experiment granting monetary awards to students who gain university admission. The authors show how the program leads to significant positive effects for females but not for males. Differences in gender-incentive interaction also emerge from the field experiment by Angrist et al. (2009). In that study, researchers randomly assigned a sample of students enrolled at a Canadian university to one of three different treatments: the first group was provided with a set of support services (e.g., extra tutoring); the second group was offered financial rewards for good academic scores; the third one was offered a combination of support services and monetary incentives depending on academic performance. The results of the experiment show that while males did not react to any of the treatments, females academic performance improved significantly when monetary incentives were provided.

While females appear to react more than males to financial incentives awarded for achieving an exogenously set target, incentive schemes based on competition may yield opposite effects.

Gneezy et al. (2003) find that males are more prone to engage in competition than females and, in general, male performance increases more than the female one when subjects are exposed to a competition in solving mazes.

In a further laboratory experiment Niederle and Vesterlund (2007) find that, when given the opportunity to choose between a piece-rate payment scheme and a tournament, men select the tournament twice more frequently than women, suggesting that women tend to avoid competition when they have the chance to do so.³ On a similar line, Datta Gupta et al. (2013) investigate how competitiveness depends on the gender of the opponent. Knowing the gender of the opponent, subjects had the possibility to implement either a tournament or a piece-rate rewarding scheme. Also in this case, significantly more males than females opted-in for the tournament. The gender of the opponent directly influenced propensity to compete: Males compete more against females than against other males while females, tend to compete more against other females than against males.

De Paola et al. (2013) report about a similar study but in a field setting. A cohort of university students were invited to undertake a midterm exam under a tournament scheme having as reward bonus points to be added to the final exam' grade. In order to mitigate self-confidence and risk aversion concerns, student were coupled in pairs of comparable ability. The gender composition of the couples represent the experimental manipulation. They find that females are

³ Despite the self-selection issue, Niederle and Vesterlund (2007) find that the average performance of females self-sorted into the tournament scheme was not different than the average score achieved by males opting for the same regime.

as likely as males to take part in the competition and to achieve good results. In this case, the gender of the competitor does not influence asymmetrically males and females student's behavior.

Using data from a natural experiment involving high school students, Azmat and Iriberry (2010) find that, even when the incentive scheme is based solely on the subject's performance, providing information on the *relative* performance induces better performances in men but not in women. Closer to our study, Czibor et al. (2014) test in the field the different outcomes generated by the two most commonly used grading practice: The absolute (i.e., criterion-referenced) and the relative (i.e., norm-referenced) grading schemes. They find that motivated male college students increase their performance when exposed to a competitive norm-referenced scheme while female students – irrespective of their motivation – are not affected by the introduction of a degree of competition for good grades.

A further vein of literature focuses on competitive behavior at young age. Jalava et al. (2014), focus on the short-term impact of non-financial incentives devised to foster sixth grade students' effort when taking standardized tests (e.g., PISA tests). They find that boys' scores are positively affected by rank-based incentives linked to test's outcome while girls' performance increases under incentive schemes involving symbolic rewards. Gneezy and Rustichini (2004), who focus on 4th grade students, find that boys improve their performance more than girls when forced to compete in a running task. Differently, Dreber et al. (2011), adopting three sorts of tasks, find no evidence of gender difference in reaction to competition when boys and girls are forced to perform in a competitive environment.

The experiments above highlight the importance of controlling for students' ability and for individual characteristics and also point out significant gender differences. In our study we control for students characteristics, including ability, and we explore the role of gender, finding that males tend to respond to incentives as predicted by our model, while females do not.

Our work presents one substantial difference in comparison to most of the previous studies, as we do not introduce costly financial incentives, but explore whether pedagogical practices can affect students' effort. In this, our work is akin to Grove and Wasserman's (2006), who studied whether the adoption of grade incentives on problem sets improves students' exam performance and found that it had a positive effect only on freshmen.

From a theoretical standpoint, our paper relates also to the literature that studies the interactions among agents under competition or cooperation. As in Itoh (1991), agents (students in our case) should help their partner when a cooperative incentive scheme is in place. However, under our assumptions,

helping always induces free-riding, thus reducing the overall level of effort exerted by the team. In principle, in our setup students not only can help each other but also sabotage their partner by providing false or misleading information. Lazear's (1989) model predicts that sabotage should be common when a competitive scheme is adopted. In our setup, even though sabotage during the experiment is possible, it should not be credible, hence it should not be part of an equilibrium strategy. To test for this, we use data from our field experiment and we check whether the frequency and the intensity of interactions between paired up students is actually affected by the alternative incentive schemes.

Similarly to Bratti et al. (2011) we propose a theoretical model of cooperation/competition among students. Our model assumes that the incentive scheme is exogenously given and predicts that in a competitive environment individual effort should be higher than in a cooperative environment. By contrast, Bratti et al. (2011) propose a model where a cooperative learning approach may successfully emerge when the class is homogeneous in terms of students' ability, and they show that – when heterogeneity is high – free riding opportunities lead to an insufficient degree of cooperation between schoolmates, which in turn decreases the overall achievement of the group.

3 The Experimental Design

The experiment involved all the undergraduate students enrolled in the Introductory Econometrics course of the Management Study major at the University of Bologna in year 2010. The course lasted 10 weeks, with a three-hour lecture per week. In addition to the regular lectures, enrolled students had the option to participate in 5 evaluated tests yielding extra points for their course grade, up to 14% of the maximum final mark. Although participation in the test was not compulsory, 131 out of 145 students applied for it. In our analysis, we excluded from the sample the applicants who missed one or more tests, hence our final sample consists of 114 students. Test marks determined the number of bonus points for the final exam. The individual bonus was equal to the average of the marks the student obtained in the five tests.⁴

⁴ Marks in the final exam range from 0 to 30. The exam is passed with a mark equal or above 18. The bonus points ranged from 0 to nearly 4.

3.1 Test Format and Procedure

Each test was computerized and consisted of five multiple-choice questions to be answered in 50 minutes and covered all topics taught in the course right up to the last lecture before the test. The five tests were designed so as to keep the degree of difficulty constant. We ran them in the computer laboratory at the School of Economics at the University of Bologna.⁵ Desks were laid out in order to prevent students from talking during the tests (see Figure 3 in the Appendix). Given the high number of students and the capacity of the lab, two consecutive sessions took place for each test on the same day, and the questions were identical across sessions. To ensure that students of the first session could not communicate with those participating in the subsequent one, we asked the students assigned to the second session to gather at the entrance of the lab 10 minutes before the start of their session. Then students of the first session exited the lab from another door placed at the other end of the room. Moreover, we did not allow students to leave the lab in case they completed their test before the end of the 50 minutes assigned.

3.2 Treatments

The mark in each test consisted of an individual component, based on the number of correct answers in the test, and a number of extra points related to the experimental condition.

Our study included a baseline condition and two treatment conditions characterized, respectively, by a competitive and by a cooperative incentive scheme. The structure of the three experimental conditions is as follows. Under all conditions, in tests 1 and 5 students worked individually. In the two treatment conditions (cooperative and competitive) students were randomly matched in pairs at the beginning of tests 2–4⁶ and had the opportunity to exchange messages with their partner via a controlled chatroom, running on their computer. In the baseline condition students always worked individually. In all conditions, part of the incentive depended on individual effort. This was the only determinant of the

⁵ The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

⁶ In case the number of students in a given test happened to be odd, one of them was chosen at random to work individually for that test, under the same incentive scheme of baseline subjects. Hence, for subjects assigned to the treatment conditions, the probability to undertake the test individually was very low and equal to the probability of an odd number of students coming to the test times the probability of being extracted as the student working individually.

Table 1: Summary of the treatments, in tests 2, 3 and 4.

Treatment	Extra points (rounds 2, 3, 4)	Messages available
Baseline	1	No
Cooperative	$1 \cdot I(s_{j,k} \geq 1.5)$	Yes
Competitive	$2 \cdot I(s_{i,k} > s_{j,k})$	Yes

whole score in all tests under the baseline, while under the competitive and cooperative treatments the total score in tests 2, 3 and 4 depended also on the partner’s performance.⁷ Table 1 summarizes the main characteristics of the experimental conditions, which are described in detail below.

Students were assigned to experimental conditions after test 1, before test 2.⁸ At the beginning of each of tests 2, 3 and 4 students assigned to the two treatment conditions were asked whether they wanted to use the chatroom to communicate with the assigned partner. This decision was taken simultaneously and paired up students could use the chatroom only if both declared to be willing to communicate. If two students chose to communicate, for each question in the test they could send one “signal” to indicate their preferred answer and/or one short text message of up to 180 characters. Interactions were anonymous, as students were not informed of the identity of their partner. In the baseline treatment no interaction between students was allowed.⁹

In each test, the value p^q of correct answer to each question q ranged between 0.3 and 1.2 points according to the question’s relative difficulty. The overall difficulty of the five tests was kept constant,¹⁰ while the topics changed

7 The decision to run the three treatments at the same time allows us to keep constant all contextual elements (such as teaching quality, cohort effects, etc.) but opens to the possibility of spillover effects between treatments. Even though we cannot rule this out, we notice that potential spillover effects would most likely mitigate the impact of the treatments on effort, hence playing against our results. In addition, no student complained about the grading system (see also page 9 on the calibration of the points awarded in each treatment) or asked to be re-assigned.

8 Hence, students did not know their actual condition when studying for test 1. This key feature of the experimental design will be exploited in Section 5 for identification.

9 Figure 6 presents a screenshot of the graphical interface of the program used for the tests. On the left-hand side of the screen students could read the question, and the multiple-choice answers. On the top-right part of the screen there was a box from which they could send messages to their partner, while on the bottom-right section of the screen they could read the messages sent to them by their partner.

10 Each test included two easy questions (0.3 points), two medium-difficulty questions (0.6 points) and one difficult question (1.2 points).

across tests. Across all treatments, the number of points $v_{i,k}$ a student could get by correctly answering the questions of test k was:

$$v_{i,k} = s_{i,k} \cdot I(s_{i,k} \geq 1.5), \quad s_{i,k} = \sum_{q=1}^5 p_{i,k}^q, \quad k = 1, \dots, 5$$

In each test, the maximum number of points $\bar{v}$ was equal to 3. This was the individual part of the mark in the test, i.e., the component which was common to all treatments.

In the competitive treatment, student i 's mark in a test was increased by 2 extra points if her score resulted to be strictly higher than her partner's. The k th test mark $\hat{v}_{i,k}$ for student i under this incentive scheme is described in eq. [1]:

$$\hat{v}_{i,k} = v_{i,k} + 2 \cdot I(s_{i,k} > s_{j,k}), \quad k = 2, 3, 4 \quad [1]$$

This provides an incentive for both paired up students to compete.

Conversely, in the cooperative treatment, student i 's score in a test was increased by 1 extra point if the partner's net score was sufficiently good, i.e., at least equal to a fixed threshold (1.5 points). The k th test mark $\hat{v}_{i,k}$ for student i under this incentive scheme is presented in eq. [2]:

$$\hat{v}_{i,k} = v_{i,k} + I(s_{j,k} \geq 1.5), \quad k = 2, 3, 4 \quad [2]$$

Finally, students in the baseline treatment received 1 extra point in tests 2, 3 and 4. This was done so that the maximum number of bonus points per pair is constant across treatments.¹¹ This element of the design was thoroughly explained to students in class. We clarified that no treatment by design systematically granted fewer bonus points. In fact, for the design to be correctly balanced, incentives in the cooperative and competitive treatment should have the same size in expectation, i.e., holding ability constant.¹²

Notice that in our design, communication is possible only in the cooperative and competitive treatments, but not in the baseline. Hence, the treatment effects we can possibly observe are determined as the outcome of the interplay between

¹¹ Given this feature of the design, potential biases caused by *John Henry* effects are ruled out by construction.

¹² In each test the probability of getting the bonus points under the competitive treatment should be half the probability of getting the bonus points under the cooperative treatment. We tested this assumption in our data and it is never rejected at 5% level. More specifically, the estimated probability of getting the bonus points in tests 2, 3 and 4 was, respectively, equal to 0.73, 0.95 and 0.95 under the cooperative treatment; and to 0.5, 0.39 and 0.44 under the competitive treatment.

communication and the incentive scheme. Because communication is essential in the cooperative and competitive treatments, the only alternative would have been to allow for communication in the baseline. However, the baseline aims at measuring the performance of students in isolation, which would not be possible, had we allowed for communication.

3.3 Timeline of the Experiment

The experiment started in February 2010 and ended in July of the same year. In the first lecture of the course, on February 25th, the full set of instructions was distributed to students and each of them had two days to decide whether to take the tests or not. At that stage, students were not explicitly informed that they were taking part in an experiment;¹³ only at the very end of the course participating students were asked to sign a consent form authorizing the treatment of the data collected during the tests.¹⁴ In this sense, our study is a “field experiment” under the terminology defined by Harrison and List (2004).

On March 1st, during a standard class, students were asked to fill out a questionnaire collecting data on some personal characteristics (age, gender, familiarity with computers, frequency of use of e-mail and chatrooms, mother and father’s level of education). We use questionnaire answers in the econometric analysis to control for individual-specific characteristics.¹⁵

On March 22nd students took the first test. Note that at this stage students had not yet been assigned to treatments, therefore the grade they obtained in the first test can be used as a measure of their effort level before being exposed to the treatment. Students were informed of the treatment to which they had been

13 We acknowledge however that subjects might have understood the purpose of the different treatments, thus suspecting that they were part of some field experiment. In any case, the *Hawthorne effect* should apply equally to all treatments; hence, it should not affect the internal or external validity of the results.

14 Data treatment and analysis was performed in compliance with the national standards on the processing of personal information (D.L. n. 196 June 30, 2003). To match the survey and experimental data we used the university numeric student identifier (*registration number*), which uniquely identifies each student. The experimental procedures were authorized by the ethics committee of the University of Bologna (*Comitato Bioetico per la Valutazione di Protocolli di Sperimentazione*). At the end of the field experiment, a consent form was administrated to all participants, in order to ask for their authorization to the treatment of the data for scientific purposes. All students involved in the experiment signed the consent form.

15 An overview of the answers to the questionnaire is provided in Section 5, and a translation of the questions is reported in Table 8 in the Appendix.

assigned only three days later, on March 25th.¹⁶ Right after the end of the test, students were informed about their own result in the first test and about the distribution of the first test score among participants. In this way we tried to convey common knowledge of the distribution of competences and ability within the population. Section 4 shows how this is relevant from a theoretical point of view.

The remaining four tests were taken approximately every 2 weeks, in April and May 2010, with the exception of the fifth which was administered 1 week after the fourth. Questions in this last test covered the whole spectrum of the topics of the course. These two design features should ensure that the students' performance in the fifth test mostly depends on the level of effort exerted during the whole duration of the course, and not just in the previous few days. After the tests, students did not receive any feedback on the performance of the other students but were only informed of their own score. Students could benefit of the bonus points gained in the tests only if they took the final exam in June or July 2010. Before the experiment started, students were informed that the bonus points would expire after the summer exam term (no bonus would be granted in late retakes).

4 The Model

This section describes the main features of the model we use to derive theoretical predictions and shape the experimental design. After a brief characterization of the theoretical setup, we will proceed to illustrate its implications in terms of expected effort under the different incentive schemes. We first describe what happens without competitive or cooperative incentives (baseline treatment); we then characterize the optimal effort under incentives to cooperation and to competition; finally, we highlight the testable predictions of the model.

16 Students taking part in our experiment were then assigned to two groups of about 65 people each, as the computer cluster can only host up to 80 students at a time. All students assigned to the competitive treatment and half of those assigned to the baseline treatment were in the first group, while all students in the cooperative treatment and the remaining students of the baseline treatment were in the second group. Using data on students assigned to the baseline treatment, we tested whether being assigned to the first or second group affected students' performance during any one of the tests (1,2,3,4,5) or their change in effort between tests 1 and 5: we could not detect any significant difference.

4.1 General Features

We assume that students' abilities are in the interval $\theta \in [0, 1]$ and are distributed according to a non-degenerate density function $f(\cdot)$. Students choose a level of effort $e_i \in [0, 1]$, which determines their score in the tests. The dis-utility of effort is $c(e)$ where $c(\cdot)$ is the same across the population.¹⁷ We further assume that $c(\cdot)$ is such that $c'(\cdot) > 0$ and $c''(\cdot) > 0$.

The expected score in test k is increasing in ability and effort¹⁸ and is given by the following expression:

$$s_{i,k} = e_{i,k} \cdot \theta_i \cdot \bar{v} \quad [3]$$

Thus, we assume that the productivity of effort is higher for higher-ability students and that only students with $\theta = 1$ can get the maximum score ($\bar{v}$) if they exert the maximum level of effort ($e_i = 1$).

We assume that students choose their level of effort twice: the first time when the course starts, before test 1, hence before being assigned to the treatments; the second time when the assignment to treatments takes place, i.e., after test 1. We denote the former by $e_{i,1}$ and the latter by $e_i = e_{i,k}$, $k = 2, \dots, 5$.

When students choose e_i , after test 1, their expected utility $U_i(B_i, c_i)$ depends on two components: (i) the cost of effort $c_i = c(e_i)$ and (ii) the bonus points to be obtained in the four remaining tests

$$B_i = \frac{1}{5} \sum_{k=2}^5 \int_0^1 \hat{v}_{i,k}(\theta_i, e_i, \theta_j, e_j) \cdot f(\theta_j) d\theta_j.$$

Note that the bonus points in the two treatment conditions are partly determined by the interaction with the partner.

4.2 Baseline Treatment

A student assigned to the baseline treatment does not interact with any other student. Remember that to get the bonus points from the tests, the student needs to get a sufficiently high score. If $s_{i,k} < 1.5$, student i obtains zero points

¹⁷ The disutility of effort can be thought both as the mental effort of being concentrated on study for a certain amount of hours and as the cost of spending those hours studying instead of meeting friends or doing some other activity.

¹⁸ In this context *ability* represents the cognitive or academic skills of the student, while the notion of *effort* refers to the willingness of a given student to spend more time on the books in order to better understand the contents of the course.

for test k . This implies that in principle there are two possibilities: either $e_i^{\text{BL}}(\theta_i) > 0$ and $v_{i,k} > 0$, $k = 2, 3, 4, 5$ Z, or $e_i^{\text{BL}}(\theta_i) = 0$ and $v_{i,k} = 0$, $k = 2, 3, 4, 5$. The latter result emerges when the student's ability is so low that the cost of the effort required for him to reach the threshold is higher than the marginal utility he would get from the additional points. Albeit theoretically possible, the latter case is practically implausible, hence we assume it away. Under this assumption, the expected number of bonus points he gets from tests 2 to 5 is

$$B_i^{\text{BL}} = \frac{1}{5} \cdot (4 \cdot e_i \cdot \theta_i \cdot \bar{v}) + \frac{3}{5} \quad [4]$$

As a consequence, his expected utility only depends on his own type and on the chosen level of effort, and it is not affected by the distribution of efforts and types among other students. Let us assume that the expected utility function is continuous and twice differentiable, it is separable in the bonus points and effort and has a single maximum in the interval $[0, 1]$. Let us denote by $e_i^{\text{BL}}(\theta_i)$ the optimal level of effort in the baseline treatment, for a student with ability θ_i .

4.3 Competitive Treatment

In order to model students' behavior under the two treatments and to derive predictions, we look for the equilibrium in the Bayesian-Nash games where students have private information about their own ability and a common knowledge of the distribution of ability in the population.

Under the competitive scheme, students get bonus points if their score is higher than their partner's. Equation [5] describes the expected number of bonus points in this case:

$$\begin{aligned} B_i^{\text{comp}} &= B_i^{\text{BL}} - \frac{3}{5} + \frac{3}{5} \cdot 2 \cdot \Pr(e_i \cdot \theta_i > e_j \cdot \theta_j) \\ &= B_i^{\text{BL}} - \frac{3}{5} + \frac{6}{5} \cdot \int_0^{\theta_i \frac{e_i}{e_j}} f(\theta_j) d\theta_j \\ &= B_i^{\text{BL}} - \frac{3}{5} + \frac{6}{5} \cdot F\left(\theta_i \cdot \frac{e_i}{e_j}\right) \end{aligned} \quad [5]$$

Under regularity assumption on the distribution of ability in the population, it can be shown that:

$$\frac{\partial B_i^{\text{comp}}}{\partial e_i} = \frac{\partial B_i^{\text{BL}}}{\partial e_i} + \frac{6}{5} f(\Phi_j(e_i)) \Phi'(e_i) \quad [6]$$

where Φ_j is the mapping from the effort to the individual ability.¹⁹ Now, given that $\Phi(e)' = 1/e(\theta)'$, we can rewrite eq. [6] as

$$\frac{\partial B_i^{\text{comp}}}{\partial e_i} = \frac{\partial B_i^{\text{BL}}}{\partial e_i} + \frac{6}{5} f(\theta_i) \cdot \frac{1}{e_i'} > \frac{\partial B_i^{\text{BL}}}{\partial e_i}$$

Since we assumed that $e_i^{\text{BL}} > 0$, then it follows that $\left. \frac{\partial U_i}{\partial B_i^{\text{BL}}} \frac{\partial B_i^{\text{BL}}}{\partial e_i} \right|_{e_i^{\text{BL}}(\theta_i)} = c'(e_i^{\text{BL}}(\theta_i))$. As a consequence, $\left. \frac{\partial U_i}{\partial B_i^{\text{comp}}} \frac{\partial B_i^{\text{comp}}}{\partial e_i} \right|_{e_i^{\text{BL}}(\theta_i)} > c'(e_i^{\text{BL}}(\theta_i))$, which implies that the optimal level of effort in the competitive treatment must be higher than in the baseline.

4.4 Cooperative Treatment

Under this scheme, each student has a clear incentive to share information (in tests 2, 3 and 4) because the mark depends also on their partner's effort. Notice that by exerting an effort equal to $e_i^{\text{BL}}(\theta_i)$ in the cooperative treatment student i can be sure to get the bonus, if he shares his knowledge with the student he is paired with. Hence, in the neighborhood of $e_i^{\text{BL}}(\theta_i)$ the expected number of bonus points from tests 2–5 becomes:

$$\begin{aligned} B_i^{\text{coop}} &= \frac{1}{5} \cdot [e_i \cdot \theta_i \cdot \bar{v}] \\ &+ \frac{1}{5} \int_0^1 3 \cdot [\bar{v} \cdot (e_i \cdot \theta_i + e_j \cdot \theta_j - e_i \cdot \theta_i \cdot e_j \cdot \theta_j) \\ &+ \mathbf{I}(e_i \cdot \theta_i + e_j \cdot \theta_j - e_i \cdot \theta_i \cdot e_j \cdot \theta_j > 0.5)] \cdot f(\theta_j) d\theta_j \end{aligned} \quad [7]$$

The first term in eq. [7] represents the points obtained in test 5, where no interaction between students was allowed, while the second term represents the bonus obtained in tests 2, 3 and 4. The assumption that information is shared by the students is crucial and it implies that the probability of answering a question correctly is given by the probability that either one of the two students knows the solution. It follows that, since $e_i^{\text{BL}}(\theta_i) > 0$:

19 In order to have a pure strategy Nash equilibria, the distribution function of ability must be non-degenerate and the mapping from ability to effort must be continuous and increasing. The requirement on the distribution of ability is plausible given the heterogeneity of the population, whereas the two assumptions on the mapping between ability and effort can be proven true, given our assumptions on the utility function. In the non-heterogeneous case, that is when the distribution of ability is degenerate, it can be easily shown that no pure-strategy equilibrium exists.

$$\left. \frac{\partial B_i^{\text{coop}}}{\partial e_i} \right|_{e_i^{\text{BL}}(\theta_i)} = \left. \frac{\partial B_i^{\text{BL}}}{\partial e_i} \right|_{e_i^{\text{BL}}(\theta_i)} - \frac{3}{5} \cdot \bar{v} \cdot \theta_i \int_0^1 \theta_j \cdot e_j \cdot f(\theta_j) d\theta_j \quad [8]$$

The second term on the right-hand side of eq. [8] is always non-positive and its absolute value increases with θ_i . It follows that $\left. \frac{\partial U_i}{\partial B_i^{\text{coop}}} \frac{\partial B_i^{\text{coop}}}{\partial e_i} \right|_{e_i^{\text{BL}}(\theta_i)} > c'(e_i^{\text{BL}}(\theta_i))$, which implies that – information being shared – each team member has an incentive to exploit the effort of the other thereby lowering his/her own effort below $e_i^{\text{BL}}(\theta_i)$.

To conclude, under the cooperative treatment, team members have an incentive to shrink their effort, and this detrimental effect of cooperation on effort is stronger for students with higher ability.

4.5 Testable Predictions

Our theoretical model predicts that, given the ability θ_i , the effort exerted by student i in the three treatments is such that:

$$e_i^{\text{coop}} < e_i^{\text{BL}} < e_i^{\text{comp}}$$

i.e., we expect that on average students randomized into the cooperative treatment exert lower effort than students randomized into the baseline treatment, whereas students randomized into the competitive treatment should exert more effort.²⁰ Conversely, in test 1, all students face the same incentives and the optimal effort depends only on their ability level, i.e., $e_{i,1}^{\text{coop}} = e_{i,1}^{\text{BL}} = e_{i,1}^{\text{comp}} = e_{i,1}$. Moreover, the model predicts that the detrimental effect of the cooperative scheme is stronger for high-ability individuals. Note that our main testable predictions involve the differential changes in effort across treatments and ability levels. Our design allows to measure those changes, as discussed in more detail in Section 5.1.

We also expect that students assigned to the cooperative treatment will use the chatroom more frequently and will use it to exchange information. Conversely, students assigned to the competitive treatment should use the chatroom less frequently and could potentially use it for acts of sabotage, i.e., to suggest the wrong answers. Note that in our setup, even though sabotage during the experiment is possible, it should not be credible, hence it should not be part of an equilibrium strategy. We collected data to verify this claim, which is discussed in Section 5.3.

²⁰ The ordering holds if the distribution of abilities is the same in the three treatments.

5 Results

In this section we first discuss the choice and the definition of our outcome measure; we then present the data and discuss the effects of the incentives to compete or to cooperate on information sharing and effort.

5.1 Measuring Effort

Our theoretical model predicts that for a given level of ability there is an ordering in the effort exerted by each student i , namely $e_i^{\text{coop}} < e_i^{\text{BL}} < e_i^{\text{comp}}$. Thus, we expect that, on average, students randomized into the cooperative treatment exert lower or equal effort than students randomized into the baseline treatment, whereas students randomized into the competitive treatment should exert more effort.

Equation [3] in our model describes the relationship between expected student's score in each test and effort, namely $s_i = \theta_i e_i \bar{v}$, where s_i is the net score of individual i , θ_i is a measure of individual ability, e_i is the effort exerted and $\bar{v}$ is the maximum score.²¹ Taking logs and allowing for noise in the way in which effort generates the score, we get:

$$y_i = \zeta_i + \eta_i + \log(\bar{v}) \quad [9]$$

where $y_i \equiv \log(s_i)$ is the log of the net score of individual i , $\zeta_i \equiv \log(e_i)$ is the log of the effort exerted, while $\eta_i = \log(\theta_i) + \varepsilon_i$ and $E[\eta_i] = \log(\theta_i)$, i.e., we assume that only the idiosyncratic component ε averages to 0 for any i , while the error η_i has a possibly non-zero mean equal to an individual-specific constant.

Given the mapping between performance in the test and effort as described in eq. [9], our experimental design provides a way to measure the change in effort under reasonable assumptions. As pointed out in Section 3, we observe students' performance in test 1 – before the assignment to the treatments – and in test 5 – after exposure to the treatments. Both these tests are taken individually under all treatments, they cover similar topics, share the same structure, have the same number of multiple choice questions and are designed to keep the difficulty constant. However, by construction, the score in the first test and the effort exerted to pass it cannot be affected by the treatments because both performance and effort are pre-determined with respect to the assignment to the different incentive schemes. Conversely, the score in the last test should reflect

²¹ To simplify the notation, we remove the subscript k indicating the current round.

changes in effort induced by the treatment. This is because, as explained in Section 3, the fifth test was administered 1 week after the fourth and the questions in this last test covered the entire spectrum of the topics of the course. This should ensure that the students' performance in the fifth test mostly depends on the level of effort exerted during the whole duration of the course, hence reflecting the effect of the treatments in tests 2–4. Indeed, moving from eq. [9] and comparing the performance in test 5 and 1, we have $y_{i,5} - y_{i,1} = \zeta_{i,5} - \zeta_{i,1} + \varepsilon_{i,5} - \varepsilon_{i,1}$. It follows that $E[y_{i,5} - y_{i,1}] = E[\zeta_{i,5} - \zeta_{i,1}]$, i.e., by looking at the change in the logarithm of score between the first and last test, we measure the change in the logarithm of effort net of the direct effect of any fixed individual-specific factor, such as ability.²²

It is important to remember that all our experimental conditions have a common individual incentive to increase effort, but they differ in the incentives to compete or cooperate. Following the theoretical predictions of our simple model, we expect an increase in effort in all treatments compared to a setup where no individual incentives are granted. Our experiment is not designed to estimate this common effect: in fact, all of our treatments have individual incentives. Instead, it is devised so as to measure the differences in the treatment effects. Our model predicts that the level of effort induced by the cooperative incentive scheme is no lower than the one induced by the baseline, while the opposite holds for the competitive incentive scheme. This ordering holds also if we consider $\log(e)$, because the logarithm is a monotonic transformation.

To test the theoretical predictions we first contrast the distribution of the change in effort under the three schemes and check for differential treatment effects over the effort distribution. We then assess the effect on the average change in $\log(e)$ by running the following regression

$$E[\zeta_{i,5} - \zeta_{i,1}] = \beta_0 + \beta_1 \text{Coop} + \beta_2 \text{Comp} \quad [10]$$

where β_0 represents the average change in $\log(e)$ under the baseline, β_1 is the average differential change in $\log(e)$ under the cooperative scheme compared to the baseline and β_2 is the average differential change in $\log(e)$ under the competitive scheme compared to the baseline. The theory predicts $\beta_1 < 0$ and $\beta_2 > 0$.

22 As explained in Section 3, in tests 2, 3 and 4 students assigned to the cooperative or competitive treatments can use the chatroom (to share information with a partner or to sabotage an opponent). Therefore, scores in tests 2, 3 and 4 are not apt to measure individual performance, given that we do not know what the performance of each member of the pair would have been, had he/she worked alone. For this reason, tests 2, 3 and 4 cannot be used to measure the change in individual effort, whereas, as already pointed out, a clean measure of effort can be obtained by comparing tests 5 and 1.

There is also the additional prediction that $\beta_0 = 0$, i.e., no change in effort under the baseline treatment. Indeed, our model does not consider the possibility that students' performance may increase across tests simply thanks to the fact that they are becoming more accustomed to the type of tests and the way these are performed in the laboratory. Such an effect, however, would be common across treatments and would not affect our theoretical predictions. If it occurs, in practice the estimate of the intercept β_0 will be positive. More generally the intercept captures every factor that influences performance and is constant across treatments.

5.2 Data and Descriptive Statistics

Among the 145 students attending the course, 131 applied for participation in the experiment. Our elaborations are based only on the records of the *stayers*, i.e., 114 students who participated in all 5 tests. We exclude from the elaborations the records of 17 students who missed at least one test, of these, 10 students assigned to the baseline treatment, 2 students assigned to the cooperative treatment and 5 students assigned to the competitive treatment (see Table 7 in the Appendix).²³

The samples of stayers are balanced across treatments with respect to observed pre-determined characteristics: we do not detect differences in the distribution of the score of the first test (score 1) and the average mark in previous exams (GPA) between any two treatments (baseline, cooperative, competitive) at any conventional level of confidence (see Table 2). Figures 4 and 5 in the Appendix report the empirical probability distribution of the pre-treatment variables (the score in the first test and the average mark in previous exams).

Table 1 also reports the mean value of several other individual characteristics – obtained from the subjects' answers to the questionnaire – and p-values of tests aimed at detecting differences in these characteristics across treatments.²⁴

In general, the overall sample is well balanced across treatments. There are some exceptions: the frequency of use of e-mail is significantly higher in the baseline treatment than in the competitive and in the cooperative treatments. Significant differences emerge also in terms of the educational level achieved by the students' fathers (but not the mothers).

²³ Six of these students were late for the third test and were thus excluded from that test as, due to technical reasons, when the experimental session starts, additional participants can be added only by shutting down and restarting the entire session. Students were informed beforehand that not being on time for the test would result in being excluded from the test session.

²⁴ We contrasted averages across treatments by means of linear and nonlinear regressions.

Table 2: Mean value of individual characteristics, by treatment and *p*-values of the test for the null of equal averages across treatments.

Characteristic	Mean				<i>p</i> -Values		
	Pooled	Baseline (BL)	Cooperative (COOP)	Competitive (COMP)	BL vs COOP	BL vs COMP	COOP vs COMP
GPA	24.8	24.8	24.9	24.8	0.808	0.883	0.928
Score 1	1.8	1.8	1.9	1.7	0.520	0.564	0.220
Age	21.7	21.6	21.9	21.5	0.280	0.736	0.161
Gender (male)	47.4%	40.5%	51.2%	50.0%	0.346	0.418	0.915
Freq. mail ^a	46.7%	62.9%	43.2%	33.3%	0.098	0.017	0.396
Freq. chat ^a	54.3%	54.3%	51.4%	57.6%	0.803	0.785	0.602
Freq. pc ^a	43.8%	45.7%	40.5%	45.5%	0.658	0.983	0.678
Father edu. ^b	31.4%	42.9%	13.5%	39.4%	0.008	0.772	0.017
Mother edu. ^b	29.5%	37.1%	24.3%	27.3%	0.241	0.386	0.778
Parental edu. ^c	40.0%	45.7%	27.0%	48.5%	0.102	0.819	0.067
Risk aversion	6.0	6.0	6.1	5.8	0.955	0.563	0.521
Trust 1	4.9	4.7	4.9	5.0	0.609	0.430	0.766
Truster (1) ^d	34.3%	22.9%	37.8%	42.4%	0.171	0.089	0.696
Trust 2	3.8	3.7	3.8	4.0	0.830	0.546	0.689
Truster (2) ^d	21.0%	17.1%	21.6%	24.2%	0.632	0.471	0.794
Observations ^e	114	37	41	36			

Notes: An individual is truster if his/her answer on the scale is higher or equal to 6. ^aBinary indicator for whether chatroom, pc or e-mail are used frequently, i.e., more than once a day. ^bBinary indicator for whether the father (or the mother) has higher education qualifications (i.e., if he/she is qualified to college level or higher). ^cBinary indicator for whether at least one parent (the father or the mother) has higher education qualification (i.e., if he/she is qualified to college level or higher). The significance of differences across treatments is estimated by means of simple linear and nonlinear regressions (logit) for binary indicators. *p*-Values are reported. ^dThe variable truster (1) [truster (2)] is a binary indicator that takes the value 1 if the answer of the individual on the trust 1 [trust 2] scale is above 6 and the value 0 otherwise. ^eSample statistics on GPA, score 1 and gender refer to 114 individuals; the remaining statistics refer to those students who answered the questionnaire, i.e., 105 students.

To detect the role of interaction effects between treatments and students’ ability, we consider the GPA: students participating in the experiment are third-year students taking exams in the last quarter of the third year; therefore, their academic history can be a reliable proxy of their academic skills. In line with the most recent empirical evidence from Italy (AlmaLaurea 2009), in our sample females tend to perform significantly better than males in terms of GPA (Females = 25.3, Males = 24.3, Wilcoxon test = *p*-value 0.0097). We classify

subjects as “high ability” if their GPA score is above the median GPA score in the sample.²⁵

5.3 Communication and Treatments

Students under both treatment schemes had two ways of communicating: they could send text messages and hints.²⁶ Messages and hints were limited in two ways. First, students could not send any useful information to identify themselves (under penalty of exclusion from the test); second, for each of the five questions asked in a test, a student could send and receive only one message of each type.

Table 10 in the Appendix reports descriptive statistics on the use of the chatroom and the giving out of hints by subjects who participated in all tests (the *stayers*) and were paired up in each test with a student who did the same.²⁷ The figures in Table 10 suggest that almost everybody under the cooperative treatment accepted to use the chatroom in tests 2, 3 and 4, resulting in a large fraction of students with an active chatroom during each test (79% in test 2, 95% in test 3, 90% in test 4). For those stayers who were always paired up with a stayer, the chatroom was active at least once in 97% of the cases and always in 74% of the cases.²⁸ These proportions are reduced by between 25% and 70% when we look at the behavior of students under the competitive treatment, and generally both the acceptance rate and the fraction of subjects whose chatroom was active during each test are significantly different across treatments at 5% level.

25 By taking the median as reference for the classification, we guarantee that the two groups are similar in size. We checked the robustness of our results to different choices of the threshold for the ability level: we considered the 75th and 66th percentile instead of the median. Results are robust to these changes. Regression results are not reported for brevity but are available from the authors upon request.

26 The hint consisted in a simple message informing the receiver that the sender believed a certain answer to be the right one. The sender could suggest a different answer with respect to the one actually selected in the test.

27 Table 9 in the Appendix presents descriptive statistics on the type of pairs: stayers could either be paired up with other stayers, or with non-stayers, or remain alone. The vast majority of pairs consists of stayers and the pattern in the use of chatroom and hints is unaffected if we do not restrict our attention to pairs of stayers. The descriptive statistics on the use of chatroom and hints on the full sample are not reported for brevity but they are available from the authors upon request.

28 At the beginning of the exam the students were asked to fill in their registration number. Then they had to choose if they wanted to use the chatroom or not. The chatroom was activated only if both members of the pair chose to use it.

Similar differences between treatments emerge if we look at the average number of messages sent by students and their length and if we look at the number of hints sent. Under both schemes, students used the chatroom more frequently than the hint but the number of messages and hints sent under the competitive scheme was significantly lower than under the cooperative scheme, where it was closer to the maximum number that could be exchanged (five messages and five hints per test).

We considered the data on the pairs in each test and compared the answers of the members: Table 3 shows that members of the pairs under the cooperative scheme tend to give the same answer much more frequently than the students under the competitive scheme. The difference is of about 25 percentage points in all tests. We also check whether there is a significant correlation between the activation of the chatroom and total pair performance (the sum of the test score of the pair’s members). We use the data of all pairs in test 2, test 3 and test 4 and present results in Table 4: we detect significant differences between the competitive and cooperative treatment in how using the communication tools affects pairs’ performance, but we cannot detect significant differences in the role of the average GPA of the members of each pair. The use of the chatroom is positively correlated with performance only in the cooperative treatment: here, the predicted total score of the pair increases from 3.769 to 5.135 points when the chatroom is active, while in the competitive treatment the variation is negligible. While these estimates do not have a causal interpretation, due to potential self-selection into information sharing, the observed pattern of information exchange across treatments can be interpreted as a positive response to the incentives: students understood the different mechanisms underlying the two different schemes and behaved accordingly as far as exchange of information was concerned.

Table 3: Fraction of pairs in which the members of the pair gave the same answer.

	Test 2	Test 3	Test 4
Cooperative	0.65	0.89	0.89
Sample size	19	19	20
Competitive	0.43	0.65	0.64
Sample size	15	16	17
Difference	0.23	0.25	0.24

Note: Only pairs where both members participated in all five tests are included. During some tests, some students who participated in all five tests may have been paired up with students who later dropped out: we have excluded these pairs from the calculation reported in the table.

Table 4: Ordinary least squares estimates of the effects of chatroom activation on a pair’s performance.

Variables	Coefficient	[Standard error]
Active chatroom	0.045	[0.378]
Coop· active chatroom	1.321**	[0.647]
Average GPA	0.426**	[0.172]
Coop· average GPA	−0.280	[0.207]
Coop	6.598	[5.163]
Constant	−6.483	[4.262]
Observations	106	
R ²	0.2	
Predicted score for average GPA 25, with no active chatroom		
Cooperative	3.769	
Competitive	4.160	
Predicted score for average GPA 25, with active chatroom		
Cooperative	5.135	
Competitive	4.205	

Note: One star, two stars, three stars for significant differences at the 10%, 5% and 1% level, respectively. We make use of 106 observations which are the different couples formed by the students in tests 2, 3 and 4 (34 couples out of 68 students in test 2, 35 couples out of 70 students in test 3, 37 couples out of 74 students in test 4). See Table 9 for more details.

Additional evidence in support of our finding that communication was used very differently in the competitive and in the cooperative treatments comes from the comparison between the hints sent and the answers given by the subjects. If we restrict our attention to the pairs where both members participated in all tests, we observe that under the cooperative treatment, the hint sent coincides with the answer given by the sender in 295 out of 323 cases (91.3%), while in the competitive in 14 out of 33 cases (42.4%). This confirms that subjects sent much fewer hints in the competitive treatment, and when they did so, they often had a deceptive intent.

We interpret the observed pattern of information exchange across treatments as a positive response to the incentives: students understood the different mechanisms underlying the two different schemes and behaved accordingly as far as exchange of information was concerned.

5.4 Treatment Effects

Figure 1 depicts the empirical distribution of the change in effort – that we measured as $\log(\text{net score } 5) - \log(\text{net score } 1)$ (see Section 5.1 for more details) – across

Figure 1: Box-plot showing the distribution of the change in effort, i.e., $\log(\text{net score } 5) - \log(\text{net score } 1)$ (see Section 5.1 for more details), across treatments.

treatments.²⁹ The vertical blue line represents the median of the distribution, the left hinge of the box indicates the 25th percentile, and the right hinge of the box indicates the 75th percentile. Visual inspection suggests that under the cooperative treatment, subjects perform more poorly than in the baseline treatment, while no sizeable differences emerge between the competitive and the baseline treatments.

Wilcoxon tests do not reject the hypothesis that the distribution of the change in effort is the same across treatments. These tests, however, are not appropriate if we want to establish an ordering across all three treatments. Thus, we also performed a Jonckheere–Terpstra test, a non-parametric test designed to detect alternatives of ordered class differences.³⁰

This test does reject the hypothesis that the change in effort is constant across treatments versus the alternative hypothesis that it is ordered across treatments according to our main theoretical prediction ($e_i^{\text{coop}} < e_i^{\text{BL}} < e_i^{\text{comp}}$) at 10%. p -Values of these tests are reported in Table 5, together with the mean level of the change in effort in each treatment condition.

²⁹ The statistics presented in this section are based on 113 observations in total, as we had to drop one subject, for whom it would have been not possible to compute our measure of effort as she scored 0 in the first test.

³⁰ The Jonckheere–Terpstra test is a non-parametric test for more than two independent samples, like the Kruskal–Wallis test. Unlike Kruskal–Wallis, Jonckheere–Terpstra tests for ordered differences between treatments and hence it requires an ordinal ranking of the test variable. For a more detailed description of the test, see Hollander and Wolfe (1999). The test is commonly used in experimental economics (Harbring and Irlenbusch 2003; Ferraro and Cummings 2007; Huck et al. 2010).

Table 5: Mean level of the change in effort, i.e., $\log(\text{net score } 5) - \log(\text{net score } 1)$ (see Section 5.1 for more details), by gender and treatment.

	Pooled	Males	Females
Mean change in effort			
Cooperative	0.500	0.377	0.628
Obs.	41	21	20
Baseline	0.583	0.452	0.677
Obs.	36	15	21
Competitive	0.570	0.680	0.459
Obs.	36	18	18
Wilcoxon tests (<i>p</i>-values)			
Baseline vs cooperative	0.313	0.135	0.948
Baseline vs competitive	0.745	0.442	0.721
Cooperative vs competitive	0.190	0.059*	0.713
Jonckheere–Terpstra tests (<i>p</i>-values)			
	0.088*	0.016**	0.624

Note: One star, two stars, three stars for significant differences at the 10%, 5% and 1% level, respectively. One female student assigned to the baseline scored 0 in test 1, hence we had to drop that observation here, as we could not compute the change in effort.

Results from previous experiments suggest that males and females may react differently to competitive incentives (see Section 2). Consistently with previous works, we find that the picture indeed changes when we split the sample by gender. Figure 2 reveals that the treatment effect is substantially different for

Figure 2: Box-plot showing the distribution of the change effort, i.e., $\log(\text{net score } 5) - \log(\text{net score } 1)$ (see Section 5.1 for more details), across treatments, by gender.

male and female subjects. The detrimental effect of the cooperative treatment on the change in effort compared to the competitive treatment only emerges for males, whereas for females no clear treatment effect arises.

One-sided Wilcoxon tests confirm that males' change in effort is significantly lower in the cooperative treatment than in the competitive treatment at the 10% level. No significant differences emerge instead when we compared the cooperative treatment with the baseline. In contrast, the same test does not reject the hypothesis of equal distribution of the change in effort between any two treatments in the female sample. To test whether the change in effort is ordered across treatments, we ran a Jonckheere–Terpstra test on the subsamples of males and females: for males, the test rejects at 5% the null hypothesis that the change in effort is not ordered across treatments, against the alternative hypothesis that the change in effort is ordered according to what is predicted by the theory; no effect is detected for females. *p*-Values of these tests are reported in Table 5.

A second noticeable difference between males and females emerges from Figure 2. For males, the variance in the change in effort is higher in the competitive treatment, as compared to the baseline treatment, while for females the opposite holds true.³¹

The result for males is in line with the findings of several previous papers (see Bull et al. 1987; Eriksson et al. 2009), but it is also in sharp contrast with what we observe for females. Previous results from the experimental literature suggest that females tend to be more risk averse than males (Croson and Gneezy 2009), and that the heterogeneity in terms of risk preferences seems to be one of the factors that explains why competitive environments tend to induce a higher variability of effort as compared to situations in which incentives are provided on a piece-rate basis (Eriksson et al. 2009). Competitive incentives boost effort for individuals who are less risk averse but decrease it for those with a stronger aversion to risk. In the presence of substantial heterogeneity in risk aversion competition may thus increase effort variance.

We checked if these arguments can explain the gender difference that we observe in our sample. Our survey measure of risk aversion (see Table 8 in the Appendix) reveals that females are (weakly) more risk averse than males and less heterogeneous than their male counterpart, however, none of these

³¹ Both differences are statistically significant according to a Levene's test (males subsample: *p*-value = 0.034; females subsample: *p*-value = 0.008). Levene's test (Levene 1960) is commonly used in the experimental literature to assess the equality of variances across different samples (Croson 2001; Eriksson et al. 2009).

differences is statistically significant at any conventional confidence level.³² Taken together these figures suggest that the introduction of competitive incentives may have had an effect that is at the same time both smaller in magnitude and more uniform among females than among males, and this would be consistent with the reduction in the variance of the change in effort for females and the increase in the variance of the change in effort for males. Having only a survey measure of risk aversion, we prefer not to draw strong conclusions on this issue, which calls for further investigation.

Our theoretical model predicts heterogeneity in the effect of the incentives schemes on effort with respect to students' ability, at least for the competitive treatment. We used linear regression models to control in a parsimonious way for individual ability and for other individual characteristics, while assessing the effects of the treatments scheme on average change in effort. We ran the analysis separately for males and females as Figure 2 suggests that they react differently to incentives.

Table 6 presents the regression results. The dependent variable in each regression is the change in effort defined at the at the beginning of this Section (and discussed in detail in Section 5.1). All regressions include controls for parental education, risk aversion and trust.³³

The table has four columns and two panels. The top panel reports estimates of regression coefficients while the bottom panel reports p -values of both two-sided and one-sided tests. The theory predicts the direction in which the null hypothesis of no effect is violated ($e_i^{\text{coop}} < e_i^{\text{BL}} < e_i^{\text{comp}}$): by exploiting this information, we increase the power of the t -test to detect significant deviations.

Each column of Table 6 corresponds to a different regression. Columns (1) and (2) correspond to regression models that do not allow for heterogeneity in the treatment effects with respect to students' ability: column (1) reports results for males, column (2) for females. In columns (3) and (4), for males and females respectively, we run regressions that include interactions between treatments and the ability indicator based on the average mark (GPA) in previous exams.

³² The median and the variance of the distribution of risk preferences are, respectively, 6 and 3.35 in the females subsample versus 7 and 3.99 in the males subsample.

³³ We control for parental education because it is weakly unbalanced across treatments, see Table 2. In addition, we control for attitudes toward risk and for trust because these characteristics may in principle be correlated with the students' behavior in the experiment. We also ran a regression with no covariates: point estimates of the main effects are qualitatively similar to those reported in columns (1) and (2) but less precise. The results are not reported for brevity, but are available from the authors upon request. Note that, given the experimental nature of our data, covariates help to improve estimates precision without changing the results substantially, as expected.

Table 6: Ordinary least squares estimates of the treatment effects on the change in effort, i.e., $\log(\text{net score } 5) - \log(\text{net score } 1)$ (see Section 5.1 for more details): benchmark specification. Males and females.

Variables	(1)	(2)	(3)	(4)
	No heterogeneity with ability		Heterogeneity with ability	
	Males	Females	Males	Females
Constant	0.268 [0.290]	0.318 [0.304]	0.137 [0.353]	0.306 [0.373]
Cooperative	-0.113 [0.183]	-0.167 [0.217]	0.113 [0.271]	-0.086 [0.363]
Competitive	0.368* [0.194]	-0.230 [0.209]	0.509* [0.280]	-0.338 [0.349]
Coop · high ability			-0.485 [0.413]	-0.226 [0.444]
Comp · high ability			-0.213 [0.410]	0.035 [0.435]
High ability			0.122 [0.288]	0.117 [0.312]
High parental education	-0.261* [0.155]	-0.165 [0.183]	-0.338* [0.178]	-0.234 [0.205]
Frequent use of e-mail			0.146 [0.166]	-0.255 [0.203]
Risk aversion	0.044 [0.038]	0.076 [0.046]	0.042 [0.039]	0.097* [0.050]
Truster (1)	0.111 [0.159]	0.045 [0.193]	0.137 [0.169]	0.053 [0.207]
Observations	50	54	50	54
R ²	0.194	0.087	0.240	0.141
p-Values for the null of no effect against two-sided or one-sided H_1				
(R) $\equiv H_1: \beta > 0$; (L) $\equiv H_1: \beta < 0$				
Competitive				
One sided (R)	0.029**	0.864	0.035**	0.834
Two sided	0.059*	0.272	0.070*	0.333
Cooperative				
One sided (L)	0.268	0.221	0.661	0.407
Two sided	0.537	0.442	0.678	0.814
Competitive for high ability				
One sided (R)			0.163	0.859
Two sided			0.326	0.282

(continued)

Table 6: (continued)

Variables	(1)	(2)	(3)	(4)
	No heterogeneity with ability		Heterogeneity with ability	
	Males	Females	Males	Females
Cooperative for high ability				
One sided (L)			0.106	0.128
Two sided			0.213	0.257

Note: High parental education is a binary indicator that takes the value 1 if the highest qualification of at least one of the parents of the individual is above high school and 0 otherwise. Standard errors in brackets. Three stars, two stars and one star for significant effect at the 1%, 5% and 10% level respectively. The sample size relevant for the regressions is 104 instead of 114, because we included control covariates and 9 students did not answer the questionnaire. Besides, we had to exclude one additional record because one female student scored 0 in test 1, hence we could not compute the log of the test score 1 for her.

High parental education is a binary indicator that takes the value 1 if the highest qualification of at least one of the parents of the individual is above high school and 0 otherwise. Standard errors in brackets. Three stars, two stars and one star for significant effect at the 1%, 5% and 10% levels, respectively. The sample size relevant for the regressions is 104 instead of 114, because we included control variates and 9 students did not answer the questionnaire. Besides, we had to exclude one additional record because one female student scored 0 in test 1, hence we could not compute the log of the test score 1 for her.

Results in Table 6 confirm previous results on the differential effects across treatments: there is evidence of a significant increase in the change in effort under the competitive compared with the baseline treatment for males but not for females. The effect for males is about 0.368 (see column (1)); it is statistically significant at 10% (see two sided p -value for the competitive treatment in column (1)) and non-negative at 5% (see one sided p -value for the competitive treatment in column (1)). When we control for ability, we find that: (i) the positive incentive for males is higher (0.509, column (3)) for low-ability individuals (statistically significant at 10% and non-negative at 5%) and decreases substantially for high-ability individuals (see the coefficient of the interaction term between competitive and ability in column (3) in Table 6); (ii) there is a negative but not significant detrimental effect of the cooperative treatment for

high-ability individuals only (see the p -value for the one-sided test of significance of “cooperative for high ability,” reported in column (3) of Table 6). However, the difference in the effects of incentives between ability groups is not significant in our sample neither for the competitive case nor for the cooperative case (see coefficients estimates of interaction terms between treatment and ability indicator in column (3) and (4) of Table 6).³⁴ In sum, the magnitude of the effect ranges between 0.368 and 0.509, i.e., between 60% and 88% of a standard deviation of the change in the logarithm effort.³⁵ This represents a substantial increase, comparable in size to the one observed by Blimpo (2014), who uses monetary incentives based on the achievement of a specified score target.³⁶ According to the estimates in the first column of Table 6, the change in effort in the baseline treatment is 31% of the effort exerted to prepare the first test, whereas the change in effort in the competitive treatment is 89% of the effort exerted to prepare the first test, meaning a 58 percentage points difference between the two treatments.

For females, no statistically significant treatment effect is found, but we acknowledge that the power we have to detect differences is much lower than the power of the same test for males.³⁷ Allowing for heterogeneity in the effects by ability does not change this qualitative findings: no significant treatment effect can be detected for females.³⁸

Notice also that the regression results evidence no significant change in effort under the baseline treatment, hence we infer that subjects’ performance did not increase across tests because they become more accustomed to the type of tests and the way these were performed in the laboratory

³⁴ While our general model gives clear predictions on the overall effect of the treatments, the model-predicted change in effort for each type is not obvious, as it depends non-trivially on the distribution of types.

³⁵ Descriptive statistics on the change in logarithm of effort: mean: 0.55; standard deviation: 0.58.

³⁶ Blimpo (2014) reports that, in his field experiment, the treatment effect on performance ranges from 0.27 to 0.34 standard deviations, depending on the treatment.

³⁷ We computed the power of a test to detect differences between any two treatments for males and females separately using the descriptive statistics (mean, standard deviations and sample size) of our sample and significance level $\alpha = 0.05$. Tables 11 and 12 in Appendix B report the results of these exercise both when we consider the effort levels not controlling for covariates and when we consider the same outcome but net of controls.

³⁸ It is worth noticing that our design is suitable to study the effect of incentives to compete or to cooperate in the case of an individual interacting with a generic partner, i.e., when facing the general population of students. Another possibility, not explored here but with relevant implications for the design of the school system, would be to quantify the effect of the same incentives when agents know the gender of their partner (De Paola et al. 2013).

(see the estimates of the coefficient associated to the constant in columns (1)–(4)).

Students' ability does not play any role in determining the increase in the change in effort in the baseline. Few regressors are relevant: risk aversion and parental background attract significant coefficients in some specifications, suggesting that individuals who are risk averse tend on average to experience larger changes in effort, while males with higher socio-economic background (here proxied by highly educated parents) tend to experience smaller changes in effort, other things equal.

Previous experiments have shown that relevant gender differences emerge in terms of risk aversion, trust and trustworthiness (see Buchan et al. 2008; Eckel and Grossman 2008; see also Croson and Gneezy 2009 for an extensive review). These factors could interact with the incentives in different ways for males and females: unfortunately, we do not have enough statistical power to detect these gender-specific differential effects in our sample.

6 Conclusions

Our study investigates how two alternative incentive schemes affect students' effort via a field experiment. The design is guided by a theoretical model meant to capture the main features of the setting considered in this study. The field experiment involved students enrolled in an undergraduate course at the University of Bologna (Italy). We randomly assigned students to either a scheme where paired up students compete to get the reward or a cooperative scheme in which students obtain a bonus if their partner performs well or a baseline treatment in which students can neither compete nor cooperate. Differently from previous studies, none of our treatments involves financial incentives: in our setup incentives are always represented by extra points for their final grade. By doing so, we provide incentives to students in "the same currency" with which they are usually rewarded.

The theoretical model predicts an ordering between effort exerted by students under the different treatments. The data from our field experiment confirm the theoretical predictions: we observe an ordering between the effort exerted by students under the different treatments, with students in the competitive treatment on average exerting more effort compared to students in the baseline and in the cooperative treatment.

We break down our results by gender and show that a significant difference emerges: Only males react to incentives to compete, while we cannot detect any significant effect for females. Cooperation does not seem to foster effort, and no gender effect emerges either.

Our experimental results suggest that non-financial incentives based on competition have the potential to promote students' similarly to pecuniary incentives (see, for instance, Blimpo 2014), but at a lower financial cost, i.e., the one of grading more tests per student). Besides this, non-financial incentives are in principle not affected by the wealth status of students: this element can in fact become a concern, due to the relatively little saliency of monetary gains for wealthy students.

In our study, competition proves to work only with males, a result which is in line with other findings in different contexts (see, for example, Gneezy and Rustichini 2004; Niederle and Vesterlund 2010), where it has been shown that males tend to be more prone to compete than females. On the other hand, this general result remains highly debated in the literature, as a series of other experimental tests failed to detect such a gender gap in the attitudes towards competition (De Paola et al. 2013; Dreber et al. 2011).

In our case, under the competitive treatment the estimated increase in effort for males is between 58 and 76 percentage points higher than in the baseline. Our estimates imply that, for example, if a student studies for three afternoons to prepare test 1, then in the baseline he/she spends about one afternoon more to prepare test 5 whereas a student under the competitive scheme will spend roughly three afternoons more. Moreover, highlighting the gender differences in the effects of incentives to compete, we complement the results in Angrist and Lavy (2009) who show that financial incentives based on absolute performance are more effective for females.³⁹

The results of our experiment suggest that introducing competition based on relative performance could induce an overall increase in students' performance. This is because this form of competition is found to increase the performance of male students, while, according to our data, on average it does not significantly affect female students (this last claim should be taken with caution, however, due to limited statistical power). More in general, since females achieve on average higher GPA scores than males, the introduction of such incentives could help to narrow the (reversed) gender gap in academic

³⁹ The treatment's effect is sizable compared to the maximum possible increase in the grade (14%). One explanation for this strong response to our incentive is that students do not only gain from the bonus points but they also derive utility from winning the competition.

performance. The introduction of a competitive grading scheme on a large scale, involving all college students, may however have perverse effects which our study cannot measure. Indeed, long term exposure to this kind of incentive mechanisms may convey to students the idea that what really matters in our society is being better than the others, even at the cost of renouncing the efficiency gains that might derive from cooperation. The study of the effects that different incentive schemes have on people's ability to achieve and sustain cooperation in unrelated tasks is an interesting avenue for future research.

Our study represents one of the first exploration of the effects of alternative non-financial incentives based on grading rules on students' effort. These results are relevant for teachers and policy makers who aim at improving the efficiency of the educational system, because they suggest that pedagogical practices – that can be implemented by faculty members – can increase students' effort. The cooperative treatment is akin to teaching practices promoting cooperation among students through a collective evaluation (e.g., group take home assignments) while the competitive treatment is akin to teaching practices promoting competition among students (e.g., through grading rules based on forced grading curves). Our results show that teaching practices belonging to the former category are likely to lower the effort exerted by students compared to teaching practices belonging to the latter category. It would be interesting to extend the inquiry to different samples, to verify whether our result holds true for students with different majors (such as literature or philosophy) who are probably less trained to optimization, and for younger students in high school and middle-high school. From a different perspective, it would also be interesting to replicate the experiment in countries characterized by a collectivist culture (Hofstede 1983, 1986; Hofstede 2011) such as China. This would allow to verify whether our result holds true in a collectivist cultural environment in which cooperative attitudes are more likely to emerge than in western societies that are mainly based on an individualistic social paradigm.

Acknowledgments: The authors thank S. Altmann, G. Brunello, G. Calzolari, M. Casari, M. Cervellati, A. Ichino, P. Kampkötter, R. Rovelli, D. Sliwka, P. Vanin and seminar participants at the University of Bologna, University of Cologne and the Monash University, at the Brucchi Luchino 2010 Workshop (Padua), at the 2011 IWAAE (Catanzaro), at the 2012 Workshop on the Social Dimension of Organizations (Budapest) and at the 2012 RES (Cambridge) for their useful comments.

Funding: The authors gratefully acknowledge the support of MIUR-PRIN 2009 project 2009MAATFS_001, of the MIUR-FIRB grant no. RBF084L83 and of the Deutsche Forschungsgemeinschaft through the research group Design & Behavior – Economic Engineering of Firms and Markets (FOR 1371).

A Appendix

A.1 Laboratory

Figure 3: The laboratory arrangement.

A.2 Additional Tables

Table 7: Descriptive statistics.

	Assigned	Stayers	Descriptive statistics–stayers		
			Predetermined controls		Score 5
			Score 1	Exams’ avg	
Baseline (control)	47	37	1.80 (0.81)	24.76 (1.8)	2.91 (0.24)
Cooperative	43	41	1.92 (0.84)	24.88 (2.3)	2.80 (0.50)
Competitive	41	36	1.69 (0.74)	24.83 (1.6)	2.69 (0.53)
Full sample	131	114	1.81 (0.80)	24.83 (1.9)	2.80 (0.45)

Score 1: score in the first test. Score 5: score in the last test. Exams’ avg: average score in previous exams. Stayers: students who participated in five experimental sessions.

In Table 8, we report the precise definition of the questionnaire data used in the analysis.

Table 8: Description of the questionnaire data.

Variable	Corresponding question	Range	Coding
Gender	<i>Gender</i>	0, 1	age = male
Age	<i>Age</i>	0–100	age in years
Freq. mail	<i>How frequently do you check your e-mail?</i>	1–5	1 = “more than once per day” 2 = “at least once per day”
Freq. pc	<i>How frequently do you use the pc to study/work?</i>	1–5	3 = “at least once per week” 4 = “less than once per week”
Freq. chat	<i>How frequently do you exchange text messages via a chatroom (msn, Facebook, Google talk, skype, etc.)?</i>	1–5	5 = “never”
Father edu.	<i>Please, indicate the education level achieved by your father</i>	1–5	1 = “junior high school” 2 = “high school”
Mother edu.	<i>Please, indicate the education level achieved by your mother</i>	1–5	3 = “bachelor” 4 = “masters” 5 = “doctorate”
Risk aversion	<i>I would describe myself as a risk-averse person.</i>	1–10	1 = “fully agree” 10 = “fully disagree”
Trust 1	<i>Do you think that most people try to take advantage of you if they got a chance or would they try to be fair?</i>	1–10	1 = “people will try to take advantage” 10 = “people will try to be fair”
Trust 2	<i>Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?</i>	1–10	1 = “you can never be too careful” 10 = “most people can be trusted”

Table 9: Descriptive statistics on pair types in the full sample of students participating in the experiment.

Cooperative and competitive treatments				
(cooperative treatment only in parentheses)				
	Test 2	Test 3	Test 4	Test 2,3,4
Both stayers	68 (38)	70 (38)	74 (40)	59 (34)
Mixed (1 stayer, 1 non-stayer)	14 (4)	8 (2)	4 (2)	2 (1)
Stayer	2 (1)	3(2)	1 (0)	0 (0)
Both non-stayers	0 (0)	0 (0)	0 (0)	0 (0)
Non-stayer	0 (0)	3 (1)	5 (1)	0 (0)

Note: “Stayers” are students participating in all five tests (77 students in the competitive and cooperative treatment). Students could be paired up only in test 2, 3 and 4 in the competitive and cooperative treatment.

Table 10: Descriptive statistics on how the students use the communication tools at their disposal, by treatment and test.

Use of the chatroom: acceptance rate (1) and fraction of subjects with active chatroom (2) (over subjects)								
	Test 2		Test 3		Test 4		Tests 2,3,4	
	(1)	(2)	(1)	(2)	(1)	(2)	Active at least once	Always active
Cooperative	0.89	0.79	0.97	0.95	0.95	0.90	0.97	0.74
Sample size	38	38	38	38	40	40	34	34
95% CI**	0.8–1	0.7–0.9	0.9–1	0.9–1	0.9–1	0.8–1	0.9–1	0.6–0.9
Competitive	0.63	0.40	0.69	0.44	0.74	0.59	0.72	0.2
Sample size	30	30	32	32	34	34	25	25
95% CI**	0.5–0.8	0.2–0.6	0.5–0.9	0.3–0.6	0.6–0.9	0.4–0.8	0.5–0.9	0–0.4
Sent hints*: unconditional (1) and conditional on subjects with active chatroom (2) (over subjects)								
	Test 2		Test 3		Test 4		Tests 2,3,4	
	(1)	(2)	(1)	(2)	(1)	(2)	Always in a pair	
							Active at least once	Always active
Cooperative	2.34	2.97	2.87	3.03	3.13	3.47	8.49	9.32
Sample size	38	30	38	36	40	36	33	25
Competitive	0.46	1.17	0.15	0.35	0.41	0.70	1.28	3.20
Sample size	30	12	32	14	34	20	18	5
Sent messages*: unconditional (1) and conditional on subjects with active chatroom (2) (over subjects)								
	Test 2		Test 3		Test 4		Tests 2,3,4	
	(1)	(2)	(1)	(2)	(1)	(2)	Always in a pair	
							Active at least once	Always active
Cooperative	2.97	3.77	3.13	3.31	3.33	3.52	9.70	10.84
Sample size	38	30	38	36	40	31	33	25
Competitive	0.57	1.42	0.56	1.07	1.18	1.43	2.83	5.6
Sample size	30	12	32	14	34	14	18	5

Table 10: (continued)

Average message length* conditional on active chatroom: characters (1) and words (2) (over subjects)										
	Test 2		Test 3		Test 4		Tests 2,3,4			
	(1)	(2)	(1)	(2)	(1)	(2)	Always in a pair			
							Active at least once		Always active	
							(1)	(2)	(1)	(2)
Cooperative	58.50	11.06	55.67	10.29	54.73	10.13	56.02	10.44	50.67	9.46
Sample size	29	29	33	33	36	36	33	33	25	25
Competitive	46.96	9.87	20.17	3.35	34.77	6.10	35.43	7.24	36.92	7.07
Sample size	9	9	5	5	17	17	13	13	4	4

Note: All students participating in all the five tests (stayers) who are matched with another student participating in all tests are included. Students who are not matched up with another student but with the pc are not included; students matched up with a non-stayer student are not included.

*Only actions, i.e., messages or hints, sent by students participating in five tests (stayers) are included in the calculations. The average length of messages is computed over the sample who sent at least one message; records with zero messages are set to missing and do not contribute to the average. The average number of messages includes records of students who sent zero messages. **95% confidence intervals based on normal approximation.

A.3 Additional Figures

Figure 4: Empirical probability distribution of score 1 by treatment.

Figure 5: Empirical probability distribution of average score in previous exams by treatment.

Figure 6 is a screenshot of the graphical interface for partial exams. The interface is divided into several sections:

- Tempo rimanente: 50 minuti** (Time remaining: 50 minutes)
- Domanda 1 [punti: 0.5]** (Question 1 [points: 0.5])
- In un modello di probabilità lineare, il coefficiente associato ad una esplicativa (continua):** (In a linear probability model, the coefficient associated with an explanatory variable (continues):)
- 1) Indica il segno dell'effetto marginale dell'esplicativa ma non la sua entità** (Indicate the sign of the marginal effect of the explanatory variable but not its magnitude)
- 2) Indica la variazione in punti percentuali della probabilità di successo corrispondente ad una variazione dell'1% dell'esplicativa considerata** (Indicate the variation in percentage points of the probability of success corresponding to a 1% variation of the explanatory variable considered)
- 3) Indica la variazione della probabilità di successo associata ad una variazione unitaria dell'esplicativa** (Indicate the variation of the probability of success associated with a unitary variation of the explanatory variable)
- risposta 1)** ☐ **risposta 2)** ☐ **risposta 3)** ☐ **non so** ☐
- Non hai ancora risposto a questa domanda. Seleziona la tua risposta e premi "Salva". Potrai comunque rivedere e cambiare la tua risposta in qualsiasi momento, prima di premere "Fine".** (You have not yet answered this question. Select your answer and press "Save". You can still review and change your answer at any time, before pressing "End".)
- MESSAGGI IN USCITA** (Outgoing Messages)
- Puoi comunicare la tua risposta all'altro membro della coppia** (You can communicate your answer to the other member of the pair)
- Per farlo, seleziona la tua risposta e clicca "Spedisci"** (To do this, select your answer and click "Send")
- a SINGLE hint per question** (a SINGLE hint per question) (a SINGLE text message (180 ch. max.))
- Puoi inviare all'altro membro della coppia un solo breve messaggio di testo [max. 180 caratteri]. Per spedire il messaggio di testo premi "Invio" sulla tastiera.** (You can send the other member of the pair a single short text message [max. 180 characters]. To send the text message, press "Enter" on the keyboard.)
- MESSAGGI RICEVUTI** (Incoming Messages)
- L'altro membro della coppia non ti ha inviato alcuna risposta.** (The other member of the pair has not sent you any answer.)
- L'altro membro della coppia non ti ha inviato alcun messaggio.** (The other member of the pair has not sent you any message.)
- Domanda 1** **Domanda 2** **Domanda 3** **Domanda 4** **Domanda 5** **Fine**

Figure 6: Screenshot of the graphical interface of partial exams.

B The Experiment Power: Discussion

We computed the power of a test to detect differences between any two treatments for males and females separately using the descriptive statistics (mean,

standard deviations and sample size) of our sample and significance level $\alpha = 0.05$. Tables 11 and 12 report the results of these exercise both when we consider the effort levels not controlling for covariates and when we consider the same outcome but net of controls.⁴⁰

Table 11: Power of the two-sided and one-sided test of mean comparison across treatments and optimal size n^* for two-sided tests with equal group sizes, power 0.8 and $\alpha = 0.05$. Males and Females. Outcome: Effort.

Null hypothesis	Males			Females		
	Power		n^*	Power		n^*
	Two sided	One sided		Two sided	One sided	
Baseline vs cooperative	0.08	0.12	658	0.06	0.08	3,078
Baseline vs competitive	0.25	0.35	84	0.24	0.35	101
Cooperative vs competitive	0.34	0.46	63	0.16	0.25	161

Table 12: Power of the two-sided and one-sided tests of mean comparison across treatments and optimal size n^* for two-sided tests with equal group sizes and power 0.8 and $\alpha = 0.05$. Males and Females. Outcome: Effort, net of controls (parental education, trust, risk aversion; treatment effects included).

Null hypothesis	Males			Females		
	Power		n^*	Power		n^*
	Two sided	One sided		Two sided	One sided	
Baseline vs cooperative	0.13	0.20	215	0.12	0.18	255
Baseline vs competitive	0.49	0.61	32	0.26	0.37	86
Cooperative vs competitive	0.64	0.75	24	0.06	0.09	1,351

Table 11 shows that we have little power to detect differences between the baseline and the cooperative treatment for both males and females, while we have more power to detect differences between the baseline and the competitive treatment for both genders, even if all values are quite low. In addition, the table reports the sample size required for a test to detect a difference of the size we observe with power 0.8 and $\alpha = 0.05$ assuming constant sample sizes across groups: most of these sizes can be hardly met within a design structured as ours.

⁴⁰ For this analysis we consider as outcome the adjusted residuals of the regressions in columns (1) and (2) of Table 6, in the paper, net of parental education, risk aversion and trust but including treatment effects.

Controlling for covariates can substantially improve the power: while our ability to detect differences between the baseline and the cooperative treatment remains low, the power to detect differences between the baseline and the competitive treatment and between the competitive and cooperative treatment is reasonably high for males (see Table 12). It remains low for females.

References

- AlmaLaurea. 2009. XI Rapporto AlmaLaurea 2009. Report. Consorzio Interuniversitario AlmaLaurea – MIUR.
- Angrist, J., D. Lang, and P. Oreopoulos. 2009. “Incentives and Services for College Achievement: Evidence from a Randomized Trial.” *American Economic Journal: Applied Economics* 1:136–63.
- Angrist, J., and V. Lavy. 2009. “The Effects of High Stakes High School Achievement Awards: Evidence from a Randomized Trial.” *American Economic Review* 99:1384–414.
- Azmat, G., and N. Iriberrí. 2010. “The Provision of Relative Performance Feedback Information: An Experimental Analysis of Performance and Happiness.” Working Papers (Universitat Pompeu Fabra. Departamento de Economa y Empresa), No. 1216.
- Blimpo, M. 2014. “Team Incentives for Education in Developing Countries: A Randomized Field Experiment in Benin.” *American Economic Journal: Applied Economics* 6 (4):90–109.
- Bratti, M., D. Checchi, and A. Filippin. 2011. “Should You Compete or Cooperate with Your Schoolmates?” *Education Economics* 19:275–89.
- Buchan, N. R., R. T. Croson, and S. Solnick. 2008. “Trust and Gender: An Examination of Behavior and Beliefs in the Investment Game.” *Journal of Economic Behavior and Organization* 68:466–76.
- Bull, C., A. Schotter, and K. Weigelt. 1987. “Tournaments and Piece Rates: An Experimental Study.” *Journal of Political Economy* 95:1–33.
- Croson, R. 2001. “Feedback in Voluntary Contribution Mechanisms: An Experiment in Team Production.” In *Research in Experimental Economics*, edited by Isaac, R. M., Published online 08 Mar 2015, Vol. 8, 85–97. Elsevier Science. [http://dx.doi.org/10.1016/S0193-2306\(01\)08005-X](http://dx.doi.org/10.1016/S0193-2306(01)08005-X).
- Croson, R., and U. Gneezy. 2009. “Gender Differences in Preferences.” *Journal of Economic Literature* 47:448–74.
- Czibor, E., S. Onderstal, R. Sloof, and M. van Praag. 2014. “Does Relative Grading Help Male Students? Evidence from a Field Experiment in the Classroom.” Technical Report. IZA DP No. 8429.
- Datta Gupta, N., A. Poulsen, and M. C. Villeval. 2013. “Gender Matching and Competitiveness: Experimental Evidence.” *Economic Inquiry* 51 (1):816–35.
- De Paola, M., F. Gioia, and V. Scoppa. 2013. “Are Females Scared of Competing with Males? Results from a Field Experiment.” Technical Report. IZA DP No. 7799.
- De Paola, M., V. Scoppa, and R. Nisticó. 2012. “Monetary Incentives and Student Achievement in a Depressed Labour Market: Results from a Randomized Experiment.” *Journal of Human Capital* 10 (3):289–98.
- Dreber, A., E. von Essen, and E. Ranehill. 2011. “Outrunning the Gender Gap – Boys and Girls Compete Equally.” *Experimental Economics* 14 (4):567–82.

- Eckel, C., and P. Grossman. 2008. "Men, Women and Risk Aversion: Experimental Evidence." In *Handbook of Experimental Economics Results*, edited by C. Plott and V. Smith, Vol. 1, 1061–73. New York: Elsevier.
- Eriksson, T., S. Teyssier, and M. C. Villeval. 2009. "Self-Selection and the Efficiency of Tournaments." *Economic Inquiry* 47:530–48.
- Ferraro, J., and R. Cummings. 2007. "Cultural Diversity, Discrimination, and Economic Outcomes: An Experimental Analysis." *Economic Inquiry* 45 (2):217–32.
- Fischbacher, U. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2):171–8.
- Fryer, R. J. 2011. "Financial Incentives and Student Achievement: Evidence from Randomized Trials." *The Quarterly Journal of Economics* 126 (4):1755–98.
- Gneezy, U., M. Niederle, and A. Rustichini. 2003. "Performance in Competitive Environments: Gender Differences." *The Quarterly Journal of Economics* 118:1049–74.
- Gneezy, U., and A. Rustichini. 2004. "Gender and Competition at a Young Age." *American Economic Review* 94:377–81.
- Grove, W. A., and T. Wasserman. 2006. "Incentives and Student Learning: A Natural Experiment with Economics Problem Sets." *AER: AEA Papers and Proceedings* 96:447–52.
- Harbring, C., and B. Irlenbusch. 2003. "An Experimental Study on Tournament Design." *Labour Economics* 45 (2):443–64.
- Harrison, G. W., and J. A. List. 2004. "Field Experiments." *Journal of Economic Literature* 42:1009–55.
- Hofstede, G. 1983. "National Cultures in Four Dimensions: A Research-Based Theory of Cultural Differences among Nations." *International Studies of Management & Organization* 13:46–74.
- Hofstede, G. 1986. "Cultural Differences in Teaching and Learning." *International Journal of Intercultural Relations* 10:301–20.
- Hofstede, G. 2011. "Dimensionalizing Cultures: The Hofstede Model in Context." *Readings in Psychology and Culture* 2:3–26.
- Hollander, M., and D. A. Wolfe. 1999. *Nonparametric Statistical Methods*. 2nd edn. New York: John Wiley and Sons.
- Huck, S., G. Lunser, and J. R. Tyran. 2010. "Consumer Networks and Firm Reputation: A First Experimental Investigation." *Economics Letters* 108:242–4.
- Itoh, H. 1991. "Incentives to Help in Multi-Agent Situations." *Econometrica* 59:611–36.
- Jalava, N., J. Schroter Joensen, and E. Pellas. 2014. "Grades and Rank: Impacts of Non-Financial Incentives on Test Performance." Technical Report. mimeo, August 2014.
- Kremer, M., E. Miguel, and R. Thornton. 2009. "Incentives to Learn." *Review of Economics and Statistics* 91:437–56.
- Lazear, E. P. 1989. "Pay Equality and Industrial Politics." *Journal of Political Economy* 97:561–80.
- Leuven, E., H. Oosterbeek, and B. van der Klaauw. 2010. "The Effect of Financial Rewards on Students' Achievements: Evidence from a Randomized Experiment." *Journal of the European Economic Association* 8:1243–65.
- Levene, H. 1960. "Robust Tests for Equality of Variances." In *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, edited by Olkin, I., 278–92. Stanford: Stanford University Press. Stanford Studies in Mathematics and Statistics.
- Niederle, M., and L. Vesterlund. 2007. "Do Women Shy Away from Competition? Do Men Compete Too Much?" *The Quarterly Journal of Economics* 122:1067–101.
- Niederle, M., and L. Vesterlund. 2010. "Explaining the Gender Gap in Math Test Scores: The Role of Competition." *Journal of Economic Perspectives* 24:129–44.