

COLLEGE COST AND TIME TO COMPLETE A DEGREE: EVIDENCE FROM TUITION DISCONTINUITIES

Pietro Garibaldi, Francesco Giavazzi, Andrea Ichino, and Enrico Rettore*

Abstract—University tuition typically remains constant throughout the years of enrollment while delayed degree completion is increasingly a problem for academic institutions around the world. Theory suggests that if continuation tuition were raised, the probability of late graduation would be reduced. Using a regression discontinuity design on data from Bocconi University in Italy, we show that a 1,000 euro increase in continuation tuition reduces the probability of late graduation by 5.2% when the benchmark probability is 80%. This decline is not associated with an increase in the dropout rate or a fall in the quality of students performance.

I. Introduction

FOR many students enrolled in academic programs around the world, it takes longer than the normal completion time to obtain a degree. Interestingly, this typically happens in contexts where college tuition does not increase (actually, it often decreases) for students who remain in a program beyond its regular end. This paper shows that these two facts—the time profile of tuition and the speed of graduation—are related and suggests that if tuition were raised after the regular end of a program, the probability of late graduation would be reduced. It also suggests that this outcome would be desirable in the presence of public subsidies to education, congestion externalities, and peer effects.

We base our empirical analysis on detailed administrative data from Bocconi University in Milan, Italy. During the period for which we have information (1992–2002), Bocconi, a private institution, offered a four-year college degree in economics. This data set is informative on the question under study not only because more than 80% of Bocconi graduates typically complete their degree in more than four years, but also because it offers a unique quasi-experimental setting to analyze the effect of tuition on the probability of completing a degree within the normal time.

Upon enrollment in each academic year, Bocconi students in our sample are assigned to one of twelve tuition levels on the basis of their family income, assessed by the university administration through the income tax declaration of the

Received for publication March 5, 2009. Revision accepted for publication October 28, 2010.

* Garibaldi: University of Torino, Collegio Carlo Alberto, CEPR, and IZA; Giavazzi: Bocconi University, MIT, NBER, and CEPR; Ichino: University of Bologna, CEPR, CESifo, and IZA; Rettore: University of Padova and IRVAPP.

We thank the administration of Bocconi University for providing the data and for answering our endless list of questions. We also thank Josh Angrist, Barbara Sianesi, and seminar participants at Bocconi University, CREST conference on the “Econometric Evaluation of Public Policies”—Paris 2005, CEPR PPS 2006, CEPR EEPPE 2006, EALE 2006, ESSLE 2006, Hebrew University, Koc University, Max Planck Institute–Rostock, NBER Education Meeting 2006, Paris School of Economics, Università di Roma Tor Vergata, and Sebanici University for helpful comments and suggestions. Our special gratitude goes to Stefano Gagliarducci, whose help was determinant in preparing the data, and to Silvia Calò for further assistance.

student’s family and through further inquiries. A regression discontinuity design (RDD) can then be used to compare students who, in terms of family income, are immediately above or below each discontinuity threshold. These two groups of students pay different tuitions to enroll but should otherwise be identical in terms of observable and unobservable characteristics determining the outcome of interest, which in our case is completing the program on time. We focus on students in the last regular year of the program, exploiting the fact that their current tuition is a good predictor of the tuition they would pay if they stayed in the program one more year. Thus, students on the two sides of a discontinuity threshold in the last regular year, albeit being identical in terms of pre-Bocconi characteristics, have paid different tuitions and should expect to keep on paying different tuitions in the following year if they do not graduate on time. While the tuition already paid is sunk and has to be paid in any case to obtain a degree, the tuition in case of late graduation can be avoided with greater effort during the last year. Using this source of identification, we show that if the official tuition assigned to a student in the last regular year were to increase by 1,000 euros, the probability of late graduation would decrease by 5.2 percentage points (with respect to an observed probability of 80%). We also show that the higher probability of graduating in time is not associated with an increase in the dropout rate or a fall in the quality of students’ performance as measured by the final graduation mark.

The paper proceeds as follows. Section II presents the available international evidence on the time to degree completion and on the time profile of tuition. Section III describes the related literature. Section IV describes the data and the institutional setting, while section V shows how a regression discontinuity design can be used to identify the causal effect of interest and discusses the robustness of our results with respect to some specific features of the framework in which our evaluation takes place. Section VI discusses whether there might be efficiency reasons suggesting that continuation tuition should be increased in academic institutions. Finally, section VII concludes.

II. Time to Degree and Time Profile of Tuition Around the World

Throughout the world, a large fraction of students remain in educational programs beyond their normal completion times, and this tendency appears to have increased in recent years.

At the undergraduate level, according to Bound, Lovenheim, and Turner (2006), time to completion of a degree has increased markedly over the past two decades. Various papers

Fn 1 and policy reports confirm these findings.¹ The problem of time to completion at the Ph.D. level in the United States is well known and has attracted considerable attention. In the representative sample collected by Hoffer and Welch (2006), the median time to obtain a Ph.D was 9 years in 1978 and increased to 10.1 years in 2003 with a similar pattern across fields.

Europe is not exempt from the problem. A survey conducted by Brunello and Winter-Ebmer (2003) on 3000 economics and business college students in ten European countries, finds that the percentage of undergraduate students expecting to complete their degree at least one year later than the required time ranges from 31.2% in Sweden and 30.8% in Italy to close to 0 in the United Kingdom and Ireland. According to Häkkinen and Uusitalo (2003), the problem of reducing time to graduation has been on the Finnish government agenda since at least 1969.

The problem is particularly serious in Italy, which offers the data used in this study. Among OECD countries this is the one with the smallest employment rate in the age bracket of 25 to 29 years, the highest enrollment rate in education in this age bracket, and the (second) lowest university graduation rate in the age bracket 35 to 44 years.² This is not because these Italian youths drop out of their schools; otherwise there would not be so many of them registered as “nonemployed, in education.” The fact is that Italian students have an abnormal tendency to extend their stay in a university program beyond the normal completion time, as documented in Dornbusch, Gentilini, and Giavazzi (2000). Ministry of Education data show that while on average, the mean legal duration of an Italian university program was 4.39 years, in a representative sample of 1995 graduates, the median effective duration was 7.00 years and the mean was 7.41, and this tendency appears to be common to all fields. Moreover, out of 1,684,993 students enrolled in Italian universities during the 1999–2000 academic year, 41.1% are classified as *Fuori Corso*, that is, their enrollment in the university system extended beyond the legal length of their program. Of the 171,086 graduates of the same year, 83.5% obtained their degree as *Fuori Corso* students.³

Fn 3 Interestingly, while throughout the world obtaining a degree within the normal completion time is becoming the exception rather than the rule, university tuition is normally structured in such a way that students pay the same amount for each year of enrollment, whether on schedule or beyond normal completion time. In some cases—one example is Italy—students pay less when they enroll as *Fuori Corso*. We are aware of only three cases that go in the opposite direction. In Germany, between 1998 and 2005, several *Länder*

introduced a tuition ranging between 500 and 900 euros for students who delayed their graduation, at a time when regular students paid no fee (see Heineck, Kifmann, & Lorenz, 2006). Similarly, the Finnish government passed a reform in 1992 aimed at reducing financial aid for students who delayed graduation (see Häkkinen & Uusitalo, 2003). In the same spirit, the Spanish system foresees that students pay for the credits they acquire by passing exams, but the cost of each credit increases if a student fails an exam and has to take it again (that is, with the number of times a student sits in an exam).

Although there is worldwide concern for the problem of increasing time to degree completion, outside of these three cases, there seems to be no evidence that academic institutions pay any attention to the possibility that the time profile of tuition and the speed of graduation might be related. In the rest of this paper, after a critical review of the related literature, we show empirically that a causal link may indeed exist, with possibly important efficiency consequences.

III. Related Literature

A small literature looks at the effect of financial incentives on time to complete an undergraduate degree, but its findings are ambiguous and typically not based on experimental evidence capable of controlling adequately for confounding factors and in particular for students' ability. Among the less recent nonexperimental studies, Bowen and Rudenstine (1992) and Ehrenberg and Mavros (1995) find evidence of an effect of financial incentives, in particular on completion rates and time to degree, while Booth and Satchell (1995) find no such evidence.

A more recent study by Häkkinen and Uusitalo (2003) evaluates a reform of the financial aid system in Finland aimed at reducing incentives to delay graduation, finding that the reform had some small effects in the desired direction. Similar in spirit to this Finnish study but with ambiguous findings is the paper Heineck et al. (2006), which evaluates the German reform of 1998 that introduced a fee for students enrolled in a university program beyond the regular completion time. Both studies, although based on the exogenous variation generated by a policy change, nevertheless have to rely on a comparison between similar students before and after the reform in order to identify the effect of a tuition increase on delayed enrollment.

Similarly plagued by the likely presence of confounding factors is the study of Groen et al. (2006) which evaluates the effect of the Graduate Education Initiative (GEI) financed by the A. W. Mellon Foundation. Over the years 1991 to 2001, this program distributed \$80 million to 51 departments in ten universities with the explicit goal of financing incentives aimed at reducing students' attrition and time to degree. By comparing these departments with a sample of similar control institutions, the study concludes that the GEI had a modest

¹ See, for example, Office of Scientific and Engineering Personnel (1990), Ehrenberg and Mavros (1995), Groen et al. (2006), Siegfried and Stock (2001), U.S. Department of Education (2003), State of Illinois (1999), U.C. Davis (2004), and Gao (2002). The situation is similar in Canada where a 2003 report of the Association of Graduate Studies indicates that “in many universities times to completion were longer than desired.”

² See OECD (2002).

³ Statistics for other years are similar.

impact on the outcomes under study, mostly reducing student attrition rather than increasing degree completion.⁴

Fn 4

A larger literature studies the effect of tuition and financial aid on college enrollment, an important question that we do not address here.⁵ Closer to our research goal are some recent papers that study, with mixed results, the effect of merit-based financial incentives on indicators of students' performance.⁶

Fn 5

Fn 6

To summarize, the mixed results of these related strands of literature may be a consequence of the more general ambiguity of the effects of monetary incentives highlighted by Gneezy and Rustichini (2000) and certainly require more research based on quasi-experimental evidence. This paper is, to the best of our knowledge, the first to provide such quasi-experimental evidence.

IV. The Institutional Framework

Bocconi is a private Italian university that offers undergraduate and graduate degrees in economics. The administrative data we use refer to a period (1992–1999) when Bocconi offered a four-year college degree, the same length of similar economics degrees offered by public universities at that time. Since then Italian universities, like as most other universities in Continental Europe, have shifted to three-year undergraduate degrees.

Although it differs in many ways from the rest of the Italian university system, which is almost entirely public, Bocconi matches national averages as far as the *Fuori Corso* problem is concerned. As in the rest of the country, the median (5.5 years) and the mean (5 years) effective time to obtain a degree are higher than the legal duration (4 years).⁷ In line with the national pattern is also the fraction of graduates who obtain a degree in more than 4 years, which is around 80%. Slightly lower than the national average is instead the fraction of *Fuori Corso* students among all students enrolled (30% against 44%), suggesting that Bocconi students prolong their studies beyond the regular length of the program as frequently

Fn 7

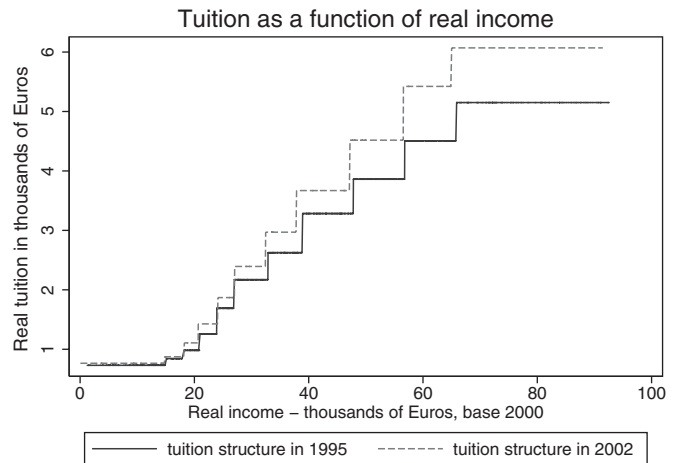
⁴ Other papers study determinants of graduation times different from financial incentives: for example, demographic characteristics in Siegfried and Stock (2001), supervisor quality in Van Ours and Ridder (2003), and labor market conditions in Brunello and Winter-Ebmer (2003). Dearden et al. (2002) study instead the effects of financial incentives on educational choices of high school graduates.

⁵ For example, van der Klaauw (2002), Kane (2003), and Dynarski (2003) and the surveys in Leslie and Brinkman (1987) and Dynarski (2002).

⁶ Angrist and Lavy (2002) find that cash awards can be very effective at increasing degree completion in low-achieving schools. Dynarski (2005) finds substantial positive effects of merit aid programs in Georgia and Arkansas on the rate of degree completion. Angrist, Lang, and Oreopoulos (2009) analyze the data of a randomized field experiment in a large Canadian university that combines "substantial merit scholarships for solid but not necessarily top, first year grades" together with or in alternative to tutoring and other auxiliary academic services. They find no effect on men but substantial effects on women. Finally, Leuven, Oosterbeek, and van der Klaauw (2006) find few or negative effects of financial rewards on measures of students' performance in the Netherlands.

⁷ It may look peculiar that the mean is smaller than the median, but the data show that at Bocconi, the majority of students who delay graduation do so for a much shorter period compared with students at state universities, so the right tail of the distribution is very short.

FIGURE 1.—TUITION STRUCTURE AT BOCCONI



Statistics for all fourth-year students who enrolled in the first year at Bocconi between 1992 and 1999. The base year for real tuition calculations is 2000.

as elsewhere but for a shorter period of time. This will be relevant for the interpretation of our results in section VI.

In the period covered by our data, students were admitted to Bocconi after an entry exam and then assigned to one of twelve tuition brackets defined in terms of family income. The highest bracket was reserved to students who accepted without discussion the highest tuition and were therefore exempt from producing their family's tax form. Since we have no income information on the students assigned to this bracket, we drop them from the analysis. Note that these students are in any case likely to be located far away from any relevant discontinuity threshold. The evolution over time of tuition in the eleven remaining brackets is described in figure 1. It should be noted that for Italian standards, tuitions at Bocconi is fairly high, ranging; for the observed eleven brackets, it is between 715 and 6,101 euros per year (at constant 2000 prices).

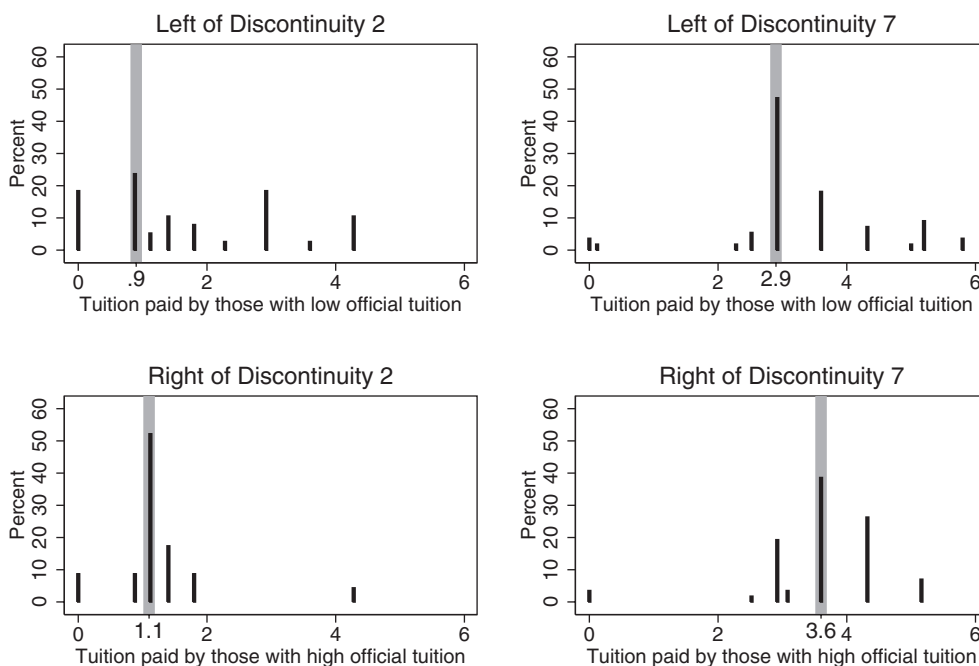
F1

A crucial feature of the admission process at Bocconi is that the university administration reserves the right to make its own reassessment of a family's ability to pay on the basis of further inquiries. As a result of this reassessment, a student may be assigned to a higher tuition level than the one implied by her declared taxable income. Moreover, for a variety of reasons (merit, orphan because of "war or similar reasons," or child of immigrants, for example), students may have a right to partial or total tuition exemption and thus pay less than what would be implied by their taxable income.

Figure 2 gives examples of the consequences of this institutional feature, using data for fourth-year students with family incomes near the second and the seventh discontinuities. Results are similar for other years and other thresholds. Starting with the top left panel, we plot the histogram of the tuition actually paid by students with family income immediately below the second discontinuity (who therefore belong to the second income bracket). These students should all pay an official tuition of 900 euros, indicated by the corresponding light bar. The dark bar of the histogram at the same level indicates that less than 25% of these students pay this official

F2

FIGURE 2.—HISTOGRAM OF OFFICIAL AND PAID TUITION FOR TWO DISCONTINUITIES IN 1998



The light bars indicate the assigned official tuitions

Histograms of the tuition actually paid by students with family income immediately to the left or to the right of the second and seventh discontinuities. The light bars indicate the official tuition that each group of students should pay. The dark bars indicate the fraction of students who pay the corresponding tuition. For example, in the top left panel, students on the left of the second discontinuity should all pay an official tuition of 900 euros, indicated by the corresponding light bar (tuition is shown in thousands). The dark bar of the histogram at the same level indicates that less than 25% of these students actually pay this official tuition. The other dark bars measure the fractions of students in this group who effectively pay other tuition levels, ranging between 0 and slightly more than 4,000 euros.

Statistics for the fourth-year students in 1998 at discontinuities 2 and 7. Results are qualitatively similar at other discontinuities and years.

tuition. The other dark bars measure the fraction of students who pay other tuitions, ranging between 0 and slightly more than 4,000 euros. The bottom left panel gives the corresponding plot for students on the right of the same discontinuity (and therefore in the third income bracket). In this case, the official tuition is higher (1,100 euros) and is paid by more than 50% of the students who should pay it in principle. The remaining students effectively pay very different tuition levels, ranging again between 0 and slightly more than 4,000 euros. The evidence in the right panels, for the seventh discontinuity, is similar. Bocconi, unfortunately, did not give us full information on the specific reasons of deviation from the official tuitions for the cases in which this happens, and thus we cannot control for it. Nevertheless, our analysis must take into account that while in the vicinity of a threshold, the assigned official tuition is binary, tuition actually paid is potentially continuous and effectively multivalued.

For all the 12,994 students enrolled in the four-year undergraduate program at Bocconi during the period 1992 to 1999 we received anonymized administrative records containing information on high school final grade and type; family income as declared to the government for tax purposes; the official tuition assigned to each student on the basis of declared family income; the tuition actually paid, which may differ from the official tuition for the reasons we explain above; the exams passed in each year and the related grades; demographic characteristics.

The support of table 1 reports some descriptive statistics suggesting that *Fuori Corso* status is correlated with indicators of lower ability and educational performance. For example, the fractions of students with top high school grades, who graduate cum laude and come from the public high school system⁸ and the top high school track,⁹ are all higher for students on time than for students *Fuori Corso*. Interestingly, the fraction of females is higher among those who graduate in time, while coming to Bocconi from outside Milan, where the university is located, does not seem to matter.¹⁰ Declared family income is on average higher for students on time, although this obviously does not say much about the causal relationship between ability to pay and *Fuori Corso* status, since family income may have a positive or negative correction with students' ability.¹¹

T1

Fn 8

Fn 9

Fn 10

Fn 11

⁸ With very few exceptions, private high schools in Italy are of a significantly lower quality, admitting students who do not survive in the public school system.

⁹ These are the only high school tracks that before 1968 granted access to university programs. In 1968 access to tertiary education was completely liberalized in Italy, so that all fields and all universities could be accessed by any student independent of his or her high school curriculum.

¹⁰ Bocconi is one of the very few Italian universities that attracts students from far away.

¹¹ Given the relatively high tuition at Bocconi, for Italian standards, students with poor family backgrounds or coming from far away with higher mobility costs, typically enroll only if they have better high school grades, which suggests higher ability.

COLLEGE COST AND TIME TO COMPLETE A DEGREE

5

TABLE 1.—DESCRIPTIVE STATISTICS BY FUORI CORSO STATUS

	Conditional on being		Of the Total
	On Time	<i>Fuori Corso</i>	
% of the 12,127 enrolled from 1992 to 1999 who:			
Are females	44.62	39.57	40.92
Are from the Milan area	40.58	40.84	40.77
Graduated from high school with top grades	28.83	22.01	23.83
Were in top high school tracks	70.40	65.98	67.16
Graduated cum laude from Bocconi	57.76	23.67	32.79
Have family income in euros	41,872	38,637	39,502
Total	26.74	73.26	100.00
% of the 10,216 fourth-year students from 1995 to 2002 who:			
Are females	47.75	40.39	41.88
Are from the Milan area	37.32	40.22	39.63
Graduated from high school with top grades	36.50	23.33	25.99
Were in top high school tracks	73.15	65.89	67.36
Graduated cum laude from Bocconi	38.10	22.08	25.31
Family income in euros	43,881	38,966	39,958
Total	20.19	79.81	100.00

Statistics for the students who enrolled in the first year at Bocconi between 1992 and 1999 and for the subset of these who were in their fourth year between 1995 and 2002.

Fn 12

In order to focus closely on the continuation decision beyond normal completion, we restrict the analysis to students in the fourth year of the program, the last regular year of studies.¹² This restriction leaves us with 10,216 students. The bottom half of table 1 reports descriptive statistics for this subsample. The comparison between the two panels of this table suggests that attrition between the first and the fourth year of the program has changed the composition of the sample in a relatively minor and expected way. To improve the comparability of treated and control subjects, the econometric analysis will be restricted to the 6,985 fourth-year students whose family income differs by no more than $\pm 3,000$ euros from a discontinuity threshold.

Note that students enrolled in the fourth, and last, regular year of the program do not know the tuition they would have to pay if they remained enrolled beyond the normal completion time. This is because they do not know with certainty the future income of their parents (family income is reassessed every year) or the future possible readjustments of the tuition structure (in terms of both levels and discontinuity thresholds) implemented by Bocconi from year to year. As a consequence, to choose their optimal level of effort during the fourth year, they must rely on a prediction of what their continuation tuition would be. Nonetheless, it is still the case that the discontinuities in the tuition system allow us to test whether students expecting higher costs of delaying graduation obtain their degree faster than otherwise identical students who expect lower costs. This is because tuition assignments are persistent in the data in the sense that official tuition in a given year is a good predictor of official tuition in the following year.¹³

Fn 13

¹² These students were observed between 1995 and 2002, since they first enrolled between 1992 and 1999.

¹³ Evidence in support of this proposition is offered by the following result. We estimate that the coefficient of a regression of the official tuition in a given year on the official tuition the year before, controlling for income and year effects, is 0.65, with a standard error of 0.009. This estimate is based on all 12,994 students enrolled at Bocconi during the period 1992 to

Tuition persistence has two implications. First, since students assigned to pay more in year t are likely to pay more also in year $t + 1$, students who pay more in their fourth year because they are just above the threshold are likely to pay more also in the event that they go *Fuori corso*. Thus, the predicted continuation tuition as a function of fourth-year income will also be discontinuous at each fourth-year tuition threshold. Second, Students who pay more in their fourth-year because they are just above the threshold are likely to have paid more also in previous years.

Does the second implication prevent us from using discontinuities in fourth-year tuition to test whether students expecting to pay more in the case of delayed tuition graduate faster? No. It is true that a student just above the threshold might have exerted more effort throughout her or his degree (and thus also in previous years) to make sure she or he graduates in four years. But this fact does not invalidate our identification since it remains true that she or he will exert more effort in his or her fourth year compared with students who happen to be just below the threshold. Could students on the right of a threshold be less likely to go *Fuori corso* because they are more able than those on the left to begin with, and therefore not because they expect to pay more in case of delayed graduation? This could happen in principle if students on the right, who expect to pay more in all years, apply only if they are very good. However, this “sorting around thresholds” based on ability can be rejected in our data, because, as we will show, we have overwhelming evidence that students on the right and on the left are identical (for example, they have the same high school grades). Therefore, the effect that we see cannot be due to differential selection at the entry margin.

To conclude, even if the quasi-experiment that we use is based on fourth-year tuition differences, the only possible

1999 for which we received the data. Thus, official tuition in a given year is indeed a good predictor of official tuition in the following year.

interpretation of what we see is that the fourth-year tuition has an effect because it proxies for continuation tuition.

V. The Evidence

A. A Regression Discontinuity Design for Our Problem

Our identification strategy is framed within the standard RDD as set by Hahn, Todd, and van der Klaauw (2001). Let y_j be the j th discontinuity point corresponding to the income level that separates tuition brackets j and $j + 1$ in the assignment rule adopted by Bocconi University. Let Y be the student's real income and τ^l be the official tuition that the student should pay according to the assignment rule, with l and h being the values of τ^l below and above the discontinuity point ($h > l$), respectively.¹⁴ Denote with τ_h^p and τ_l^p the potential treatment values—the tuition that a student in a neighborhood of the discontinuity would actually pay if the official tuition assigned to her or him were h or l , respectively. As explained in section IV, both τ_h^p and τ_l^p are in principle continuous, effectively multivalued, and possibly different from h and l , respectively. Let F_h and F_l be the potential binary *Fuori Corso* outcomes of a student under the official tuition assignment h and l , respectively. Finally, let $\tau^p = I(\tau^l = h)\tau_h^p + I(\tau^l = l)\tau_l^p$ be the observed tuition actually paid and $F = I(\tau^l = h)F_h + I(\tau^l = l)F_l$ be the observed *Fuori Corso* status, where $I(\cdot)$ is the indicator function.

Under the regularity conditions set by Hahn et al. (2001), the average effect of being assigned to the higher official tuition bracket $\tau^l = h$ (instead of the lower one $\tau^l = l$) on the observed tuition actually paid τ^p and on the observed *Fuori Corso* outcome F for a student in a neighborhood of the cut-off point are

$$E\{\tau^p|y_j^+\} - E\{\tau^p|y_j^-\}, \tag{1}$$

$$E\{F|y_j^+\} - E\{F|y_j^-\}. \tag{2}$$

These are the so called intention-to-treat effects. In order to keep the notation simple, here and below, we omit time subscripts, but in our context, these expressions identify causal effects only conditioning on time periods. This is because the composition of the pool of Bocconi students changed over the years with respect to some observables relevant to the outcome. It is therefore necessary to condition on the time period to make the students just above the cut-off point comparable to those just below it with respect to such observables.

More problematic in our context is the conversion of the intention-to-treat effects into a meaningful causal effect of τ^p on F . Following Angrist, Graddy, and Imbens (2000), the exclusion restriction, requiring that the official tuition τ^l affects the *Fuori Corso* status F only through the tuition effectively paid τ^p , needs to be satisfied. This is a plausible

restriction in our context. More critical instead is the monotonicity condition, asserting that no one is induced to pay a lower (higher) actual tuition if exogenously moved, in terms of official tuition, from l to h (from h to l). The evidence in the next section and the formal test that we perform in section VC show that this condition fails in our case. Therefore, in contrast to standard analysis, the ratio

$$\Lambda(y_j) = \frac{E\{F|y_j^+\} - E\{F|y_j^-\}}{E\{\tau^p|y_j^+\} - E\{\tau^p|y_j^-\}}, \tag{3}$$

does not identify, in our context, the average effect of a unit change in τ^p on the probability of going *Fuori Corso* at $Y = y_j$ for those who are induced to pay a higher actual tuition because their official tuition increases from l to h . Thus we present and focus only on intention-to-treat effects.

B. Graphical Evidence

Figure 3 plots nonparametric regressions of the variables τ^l , τ^p , and F on Y , respectively, for fourth-year students at the discontinuity thresholds 2 and 7, which are representative of what we obtain in the other cases. The regressions are estimated separately above and below the cut-off points so that the possible jump at the threshold may show up if it exists. Thus, these plots offer a visual image of the intention-to-treat effects defined in equations (1) and (2).

The tuition τ^p effectively paid by the student is uniformly not lower than the official tuition τ^l on both sides of the threshold. However, while at cut-off point 7, the mean value of τ^p above the threshold is higher than its mean value below it, the reverse happens at cut-off point 2. This suggests the possibility that the monotonicity condition is violated, a problem that we formally address in the section VC.

As for the main outcome of interest, the probability of observing $F = 1$ is higher below the cut-off point for discontinuity 7, but the opposite happens at the second discontinuity. Nevertheless, the mean impact of τ^p on F , which is the ratio between the jump of $Pr(F = 1)$ and the jump of τ^p , turns out to be negative at both discontinuities. This implies that in both cases, the probability of going *Fuori Corso* changes in the opposite direction with respect to the tuition effectively paid when the threshold is crossed. As we will see in the next section, however, the failure of the monotonicity condition prevents us from interpreting this ratio as the causal effect of τ^p on F .

The identification of ITT effects alone does not require the assumption of monotonicity but does require the continuity of unobservables around the thresholds. To gather evidence on the validity of these continuity conditions, we implement an overidentification test following Lee (2006). Consider the set of preintervention outcomes that meet the following two conditions: they should not be affected by the tuition system of fourth-year students at Bocconi University, but they should depend on the same unobservables (such as ability) likely to affect *Fuori Corso* status F . Two preintervention

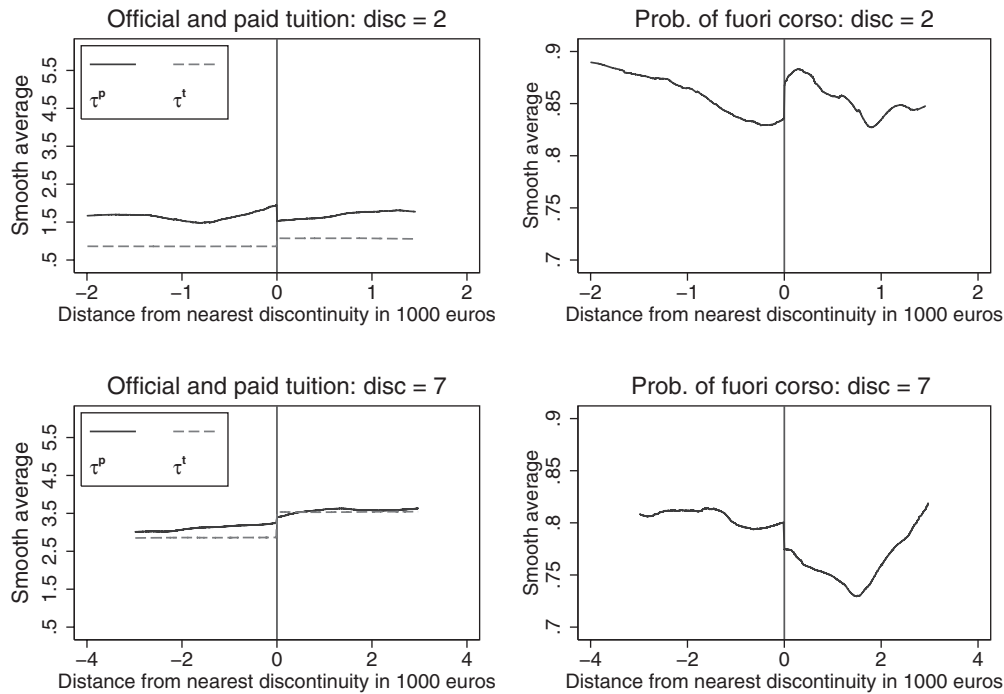
Fn 14

F3

¹⁴ In principle, a subscript j should be attached to the values of the official tuition, but since in this section, we consider only one generic threshold j , we omit this subscript to simplify notation.

COLLEGE COST AND TIME TO COMPLETE A DEGREE

FIGURE 3.—INTENTION-TO-TREAT EFFECTS



Statistics for the fourth-year students who enrolled in the first year at Bocconi between 1992 and 1999.

outcomes satisfying these requirements are family income before enrollment at Bocconi and the grade that a student receives in her or his final exam at the end of high school. Both variables are observed at least three years before the fourth year at Bocconi in which our quasi-experiment is framed. If we found that students on the two sides of a discontinuity point differed with respect to these variables, we would have to conclude that our identification strategy fails since students assigned to $\tau^t = h$ are presumably not comparable to students assigned to $\tau^t = l$ with respect to unobservables relevant for the outcome F . Figure 4 shows that no significant discontinuity of this kind emerges at the representative discontinuities 2 and 7.

F4

A formal test confirming this evidence is described in the next section, where we go beyond the graphical evidence presented so far, showing how the estimates obtained separately at each threshold can be aggregated in a single encompassing estimate.

C. Aggregation of the Mean Effects at Different Thresholds

By constructing an aggregate estimate (across all thresholds) of the average causal effect of official tuition on the probability of going *Fuori Corso*, we gain precision at the cost of losing information on how the mean effect of interest varies with Y . Following Angrist and Lavy (1999), such an estimate of the ITT parameter of interest can be obtained from the following equation

$$F = \beta\tau^t + \gamma_a + g(Y) + \delta X + \epsilon, \tag{4}$$

where $g(Y)$ is a fourth-order polynomial in Y and X is a vector of pretreatment characteristics of students. For the reasons explained at the end of section VA, we also include academic-year-specific effects γ_a in this equation. Note that in this regression, the inclusion or exclusion of students' characteristics X , observed before entrance at Bocconi, should not affect the estimate of β if, at each threshold, the assignment to treatment (high or low tuition) is orthogonal with respect to pretreatment characteristics.

Estimates of β can be interpreted as estimates of the intention-to-treat effect of assigned official tuition on the *Fuori Corso* outcome. A similar equation can be used to estimate the intention-to-treat effect on the tuition actually paid by the student. These estimates are described in table 2. As anticipated above, to improve the comparability of treated and control subjects, the analysis is restricted to observations within a window of at most $\pm 3,000$ euros with respect to each threshold.¹⁵

T2

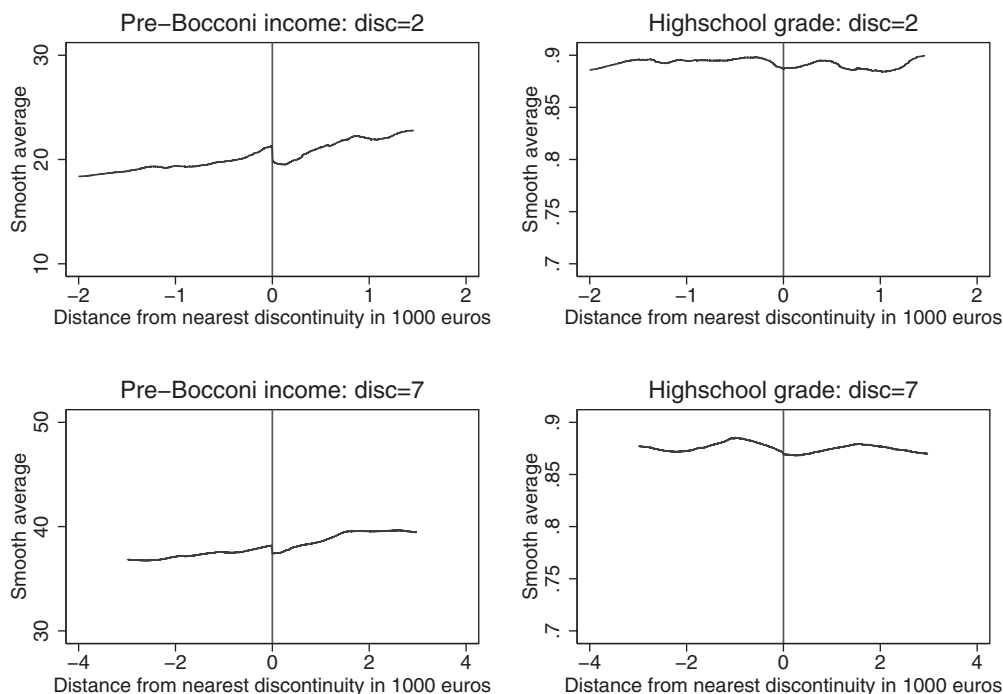
Fn 15

The intention-to-treat effect of τ^t on τ^p is reported in the first panel of the table and indicates that each additional euro of official tuition converts into .528 euro of tuition actually paid (with a standard error of .055). This is because, in the data, the downward readjustment for students on the right of a threshold is on average more frequent or larger

¹⁵Note that the sample size is smaller than the one of the full sample described in table 1. The estimates become slightly smaller in absolute size but remain statistically significant at conventional levels when other window sizes (up to $\pm 1,000$) are used. Results are available on request. As an alternative, we have also aggregated the estimates at the ten thresholds by weighting them with the inverse of their sampling variance. Results are very close to those we report.

THE REVIEW OF ECONOMICS AND STATISTICS

FIGURE 4.—EVIDENCE ON SORTING AND CONTINUITY CONDITIONS



Statistics for the fourth-year students who enrolled in the first year at Bocconi between 1992 and 1999.

TABLE 2.—REGRESSION DISCONTINUITY ESTIMATES OF THE EFFECTS OF OFFICIAL TUITION

	Paid Tuition 1	Fuori Corso 2	Paid Tuition 3	Fuori Corso 4	Paid Tuition 5	Fuori Corso 6
Official tuition	0.528 (0.055)	-0.052 (0.023)	0.531 (0.055)	-0.054 (0.023)	0.562 (0.060)	-0.047 (0.025)
Female			0.010 (0.029)	-0.031 (0.010)	0.008 (0.029)	-0.031 (0.010)
Family of origin outside Milan			-0.003 (0.028)	0.029 (0.010)	-0.002 (0.028)	0.029 (0.010)
High school grade			-1.564 (0.136)	-0.660 (0.045)	-1.564 (0.137)	-0.662 (0.045)
High school type			0.071 (0.031)	-0.054 (0.010)	0.071 (0.031)	-0.054 (0.010)
Income before Bocconi			0.008 (0.001)	-0.0008 (0.0004)	0.008 (0.001)	-0.0008 (0.0004)
Constant	2.571 (0.528)	0.870 (0.161)	3.844 (0.539)	1.500 (0.165)	13.985 (43.918)	12.220 (11.480)
Academic year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Same $g(Y)$ for all thresholds	Yes	Yes	Yes	Yes	No	No
Different $g(Y)$ for low, Medium, and high thresholds	No	No	No	No	Yes	Yes
Observations	6,985	6,985	6,790	6,790	6,790	6,790
R^2	0.529	0.0371	0.545	0.0695	0.545	0.0696

Each column reports coefficients (and robust standard errors in parentheses) estimated using regressions of the form $S = \beta\tau' + \gamma_a + g(Y) + \delta X + \epsilon$, where S is the outcome indicated in the corresponding column; τ' is the official tuition; γ_a are academic year dummies; $g(Y)$ is a fourth-degree polynomial of family income Y and X are the pretreatment characteristics included in the specifications of columns 3 to 6. In columns 5 and 6, the polynomial is allowed to differ between observations associated with three groups of thresholds: the first group is for thresholds 1, 2, and 3; the second for thresholds 3, 4, 5, and 7; the third for thresholds 8, 9, and 10.

The number of observations is smaller than the one of the full sample described in table 1 because here we restrict the analysis to students whose family income is within a window of $\pm 3,000$ euros from the closest threshold. High school type is a dummy equal to 1 for students who were in high school tracks designed to prepare for university studies (*Liceo*). The high school grade is a variable ranging between 0.6 (passing grade) and 1 (maximum grade).

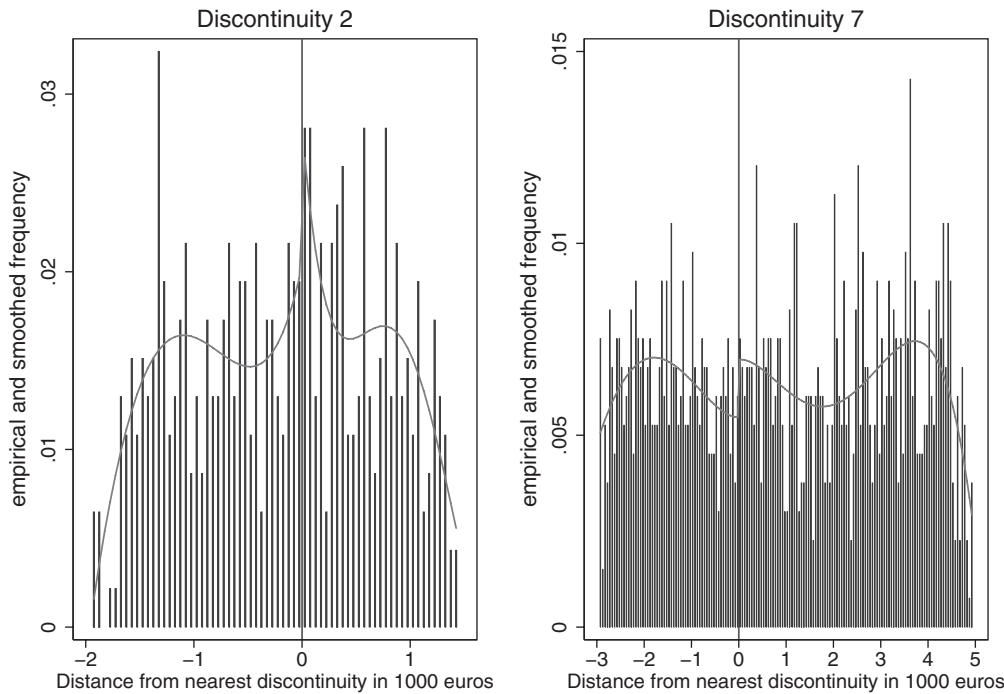
Statistics for the fourth-year students who enrolled in the first year at Bocconi between 1992 and 1999.

than the upward readjustment for students on the left. However, despite this dilution, the intention-to-treat effect of τ' on F in the second column of table 2 suggests that a 1,000 euro increase of the official tuition would decrease by 5.2

percentage points the probability of going *Fuori Corso*, with respect to a sample average of approximately 80%, with a standard error of 2.3 percentage points. The third and fourth columns of table 2 also add to the specification five

COLLEGE COST AND TIME TO COMPLETE A DEGREE

FIGURE 5.—HISTOGRAM OF FAMILY INCOME AROUND THRESHOLDS



Statistics for the fourth-year students who enrolled in the first year at Bocconi between 1992 and 1999.

characteristics of students measured before their admission to Bocconi and thus not affected by tuition assignment. Most of these pretreatment characteristics appear to affect in a sizable and statistically significant way the tuition paid by students and their probability to go *Fuori Corso*. However, the inclusion of these variables does not change in any relevant way the effect of official tuition on both outcomes with respect to the first two columns of the table.

In other words, pretreatment observable characteristics are perfectly balanced on the two sides of each threshold, which gives credit to the validity of the continuity conditions that are needed for the identification of the intention-to-treat effects in our setting.¹⁶

Particularly important for the validity of our setting is the finding that the inclusion of family income before enrollment at Bocconi is irrelevant for the estimation of the ITT of interest. This is reassuring because in principle, families can alter their declared taxable income in order to be assigned to a lower bracket. If this happened, it would result in an endogenous sorting of students around the thresholds, which would generate discontinuities in the density function at the thresholds and specifically a concentration of probability mass

immediately below them. As a result, pre-Bocconi family income would not be balanced around the thresholds, and its inclusion in equation (4) would affect the estimate of the ITTs. Table 2 suggests that this is unlikely. More direct evidence on the absence of such discontinuities is offered by the test proposed in McCrary (2008), that we have adapted to our setting: the *t*-statistics of the tests associated to the ten discontinuities are all largely insignificant.¹⁷ In figure 5, we plot the histogram of family income for fourth-year students around two representative discontinuity thresholds, the second and the seventh, and the associated estimate of the density function obtained by smoothing the histogram by a fourth-degree polynomial separately on the left and on the right of the threshold. It is evident from the figure that no discontinuity emerges at these two thresholds (as well as at the others not reported to save space). If anything, the probability mass is concentrated above the discontinuity.

To further assess the robustness of our estimates, the last two columns of table 2 allow a more flexible specification of the polynomial of family income $g(Y)$. Instead of imposing that the parameters of this polynomial are the same at all income levels, we let them be free to differ among observations belonging to three groups of thresholds, at low income (thresholds 1 to 3), medium income (thresholds 4 to 7), and high income (thresholds 8 to 10). Also in these more flexible specifications, the size and significance of the estimated ITTs is essentially not affected.

¹⁶ It is interesting to observe that for most of these variables, there is no obvious prior on the sign of the estimated effects, but no result is implausible. Females and students coming from families resident in Milan are less likely to go *Fuori Corso*, while both variables have no effect on the tuition actually paid. A higher high school grade, a high school curriculum designed to prepare for university studies (the high school type dummy), and a higher family income before arriving at Bocconi all reduce the probability of *Fuori Corso*. A higher high school grade is associated with a lower paid tuition, while the remaining two variables have the opposite effect on this outcome.

¹⁷ The *t*-values at the ten cut-off points are, respectively 0.30, 0.70, 0.9, 1.1, 0.30, -0.41, 1.29, -0.38, -0.20, -0.62.

Fn 16

Fn 17
F5

We can therefore conclude that the fourth-year official tuition has a sizable and statistically significant effect on the speed of graduation. Since fourth-year tuition is sunk, this may appear counterintuitive. But as we explained at the end of section IV, since fourth-year students do not know the tuition they would pay if they go *Fuori Corso*, this evidence suggests that they use the fourth-year tuition to predict what their continuation tuition might be. So even if what we estimate is just the causal effect of the fourth-year tuition, the fact that it is positive and statistically significant indicates that students use their fourth-year tuition to predict their continuation tuition and that the latter increases their speed to graduation.

We next explore formally whether more can be made out of our experiment and, in particular, if the IV estimand (3), the ratio of the two ITT effects, can also be given a causal interpretation. The graphical evidence presented above suggests that the possibility of this interpretation is jeopardized by the reassessment of families' ability to pay operated by Bocconi University, which implies that the official tuition initially assigned to each student does not correspond exactly to the tuition actually paid. The condition that in our setting may prevent this interpretation is, of course, not simply the lack of compliance, but, more critical, the lack of monotonicity. What we need for identification is that at each threshold, students assigned to the lower official tuition do not effectively pay more than if they had been assigned to the higher official tuition of the same threshold. Consider a student with a family income immediately below a threshold. Bocconi has a stronger incentive to open her file and reassess her income than if the student had been located immediately above the threshold, because in the first case, a small reassessment would be enough to increase the tuition this student pays. However, once the file is open, the reassessment may be large and imply a large increase in tuition. As a result, it is possible that the same student pays effectively more if assigned immediately below a threshold than if assigned immediately above, and this would imply a violation of monotonicity. A similar reasoning holds for the case of a student assigned immediately above a threshold. In this case, she will have a stronger incentive to ask for a tuition exemption than if she had been assigned by family income to a threshold immediately below.

As already noted in section VB, an indication that the problem might exist in our case is offered by the fact that at the second discontinuity threshold, the mean actual tuition paid by students assigned to the lower bracket $\tau^l = l$ exceeds the mean actual tuition paid by students assigned to the higher bracket $\tau^l = h$ (see figure 3). Similar evidence can be found at some other thresholds.

A formal test for the occurrence of defiance has been proposed by Angrist and Imbens (1995). The monotonicity condition in our case asserts that $\tau_h^p \geq \tau_l^p$, with the strict inequality holding at least for some students. In words, no one would be induced to pay a lower actual tuition if her official tuition shifted from low to high, while at least one student should be induced to pay a higher tuition in this event. This

condition is not directly testable since the two potential outcomes τ_h^p and τ_l^p of a specific student are not simultaneously observable. However, a testable implication of the inequality is that the cumulative distribution function (cdf) for those in a right neighborhood of the cut-off point should not be above the cdf for those in a left neighborhood of it at any value of its support. In our case, this implication is violated at some cut-off points. In figure 6 we present the estimated difference between the cdf on the left and the corresponding cdf on the right at the second and the seventh discontinuities (.95 confidence intervals are plotted). It is clear that the stochastic dominance hypothesis is rejected at these thresholds, suggesting that defiance occurs at least here.¹⁸

We thus conclude that the only causal effects that can be identified in our data are the intention-to-treat effects described in table 2 and discussed at the beginning of this section.¹⁹

D. Collateral Effects

It could be argued that in order to interpret these findings and draw policy conclusions, one should know whether a higher tuition makes it more likely that students drop out and whether those students who try to graduate on time do so at the expense of the quality of the learning process. In this section, we describe evidence that rejects both of these hypotheses.²⁰

To test the first one, we estimate an equation like (4) in which the dependent variable is a dummy taking a value of 1 if the student drops out after the fourth year. A 1,000 euro increase in the official tuition reduces the probability of dropping out by 0.4 to 0.6 percentage points, depending on which of the three specifications described in table 2 we use to estimate equation (4). The estimated effect is not only small but also largely statistically insignificant: there is no evidence that students assigned to a higher official tuition are more likely to drop out.²¹

¹⁸To control for year-specific effects at each discontinuity point, we estimated the difference among the two cdfs and their standard errors separately for each calendar year. Then we evaluated the weighted mean of such year-specific differences using as weights the inverse of the sampling variances.

¹⁹The working paper version of this article (see Garibaldi et al., 2007) contains a model that provides restrictive assumptions under which, despite defiance, the IV estimand identifies, for compliers, the causal effect of the tuition actually paid on the probability of going *Fuori Corso*—the LATE. The evidence in favor of the validity of such restrictive assumption is, however, weak and insufficient to justify inclusion in the published version of our research project.

²⁰To save space, we summarize results in the text. Further details, if necessary, are available from the authors.

²¹This result differs from the evidence of Dynarsky (2005) who exploits the introduction of two large merit scholarship programs in Georgia and Arkansas to show that a reduction of college costs increases significantly the probability of completing a degree. The difference between our and her findings, concerning the effect of college costs on dropout rates, may be explained by the fact that the two studies are based on different quasi-experimental situations and identification assumptions. In particular, her study focuses on tuition differences based on merit (a minimum GPA in high school and in college), while in our case, tuition differences are independent of merit.

F6

Fn 18

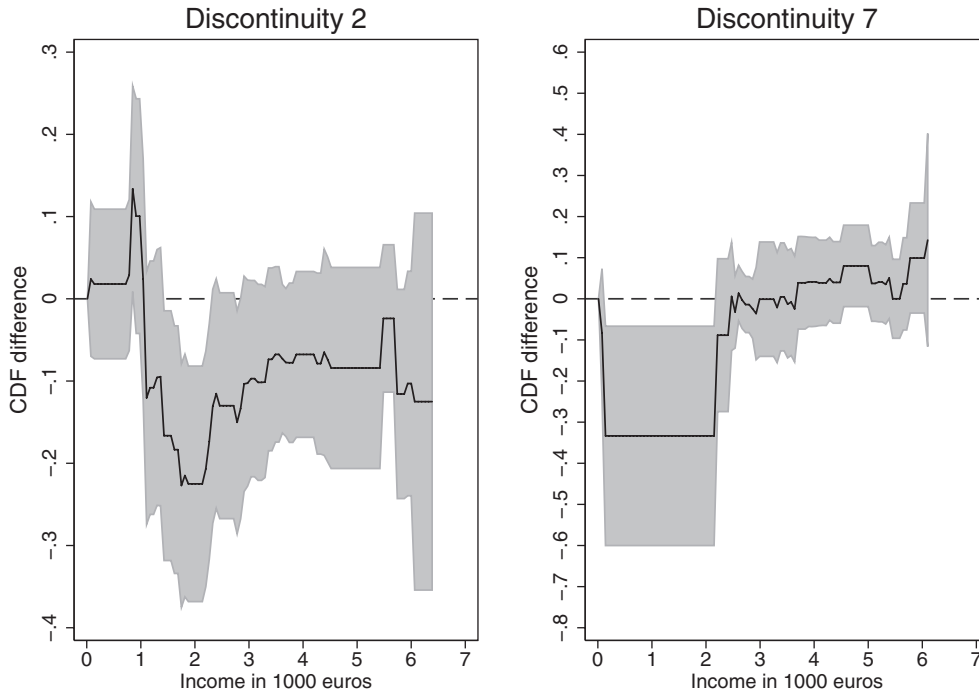
Fn 19

Fn 20

Fn 21

COLLEGE COST AND TIME TO COMPLETE A DEGREE

FIGURE 6.—A TEST OF MONOTONICITY: CDF CROSSING



For each discontinuity, the figure plots the estimated difference between the cdf of the tuition actually paid by students in a left neighborhood of the cut-off point and the corresponding cdf paid by students in the right neighborhood. The shaded area describes the 0.95 confidence intervals. The left (right) neighborhood is defined by selecting students whose family income is below (above) the cut-off point by no more than 500 euros. Statistics for the fourth-year students who enrolled in the first year at Bocconi between 1992 and 1999.

To test the second hypothesis, we estimate again an equation like (4) in which the dependent variable is the final graduation mark received by the fourth-year students in our sample who had already graduated by the time we obtained the data from Bocconi.²² This final graduation mark in principle ranges between 66 (passing level) and 110 plus honors (*Laude*), and it is determined by a committee of faculty members on the basis of the grades obtained in all the exams of the four years and in the final dissertation. In our sample, this final mark ranges effectively between 77 and honors with a standard deviation of 7 points.²³ In this case, an increase of 1,000 euros in the official tuition assigned to a student reduces her final grade only by 0.47 to 0.67 points, depending again on which specification is used among those described in table 2. Thus, if a higher tuition induces students to speed up their course work in order to finish earlier, this does not happen at the expense of the quality of the learning process inasmuch as this is measured by the final grade.

VI. Discussion

The empirical analysis has established that an increase in continuation tuition decreases the probability of late graduation. In other words, students who expect to pay more in case of delayed graduation because they are exogenously assigned to a higher official tuition seem to exert more effort

and increase their speed to graduation. The analysis has also shown that the increase in speed to graduation does not induce an increase in dropouts and does not significantly affect the quality of students' performance, at least as measured by the final graduation mark.

The size of the effect we have estimated—a 1,000 euros increase in the official tuition reduces the probability of late graduation by 5.2 percentage points, in a context in which late graduation occurs for approximately 80% of students—may at first look too large. Postponing graduation in terms of forgone income is very costly and at least one order of magnitude larger than the 1,000 euros of additional official tuition.²⁴ What we have estimated, however, is a marginal effect. The expected forgone income from delaying graduation by one year determines the speed at which students graduate given the existing tuition profile. What we find is that 1,000 euros make a significant difference at the margin once the effect of the expected forgone income is already taken into account.

How general is our result that time to degree is affected by tuition? As in any other experimental or quasi-experimental

²⁴ Around the time our data were collected, Bocconi students earned on average 25,000 euros (at 2001 prices) one year after graduation, and most of them found a job in few months. Filippin and Ichino (2005) compare data on a sample of Bocconi graduates with similar data on graduates from the State University of Milan studied by Checchi (2002). Their most conservative estimate suggests that in 2001, Bocconi graduates who had first enrolled in 1997 earned at least 1.5 times more than state university graduates of the same year. And 92% of Bocconi graduates had found a job within one year, while the same happened for only 46% of the graduates from the other institution.

²² By 2004, 1,010 students had not graduated yet.

²³ We consider honors as an additional point.

Fn 22

Fn 23

Fn 24

setting, extrapolation is problematic. Our estimates have been obtained in the context of an institution with a particular relationship between actual and assigned tuition and can thus be generalized only to institutions with a similar relationship. For instance, the effects would be different in a university that applies a hard rule, in the sense that assigned tuition must always be followed. This is a limitation of our analysis. But nevertheless our finding that time to degree is affected by tuition remains valid, even if the exact quantitative effect that we estimate cannot be generalized across institutions.

Our finding—that the speed at which students decide to learn is affected by the tuition they pay—does not necessarily mean that it is socially optimal to increase continuation tuition. We do not know much about the optimal length of the learning period for a given amount to be learned, an issue rarely explored in the literature.²⁵ Each student could choose the speed that she or he considers optimal for herself or himself, and different individual characteristics (including different preferences for work and leisure) could result in quite different “optimal” learning speeds. To make a normative argument, we need to point to reasons that individual decisions might be suboptimal. We see at least three reasons that this might happen.

The most obvious one is that students, even in some private universities, are often subsidized by the state. If students (or their families) fail to pay the marginal technological cost of their education, they will not internalize the cost to society of keeping them one more year in school and will make decisions that are socially suboptimal. Using the tuition profile to affect their incentives can then improve society’s welfare.²⁶

Another example is the evidence of peer effects in education. Peer effects in school are at work whenever there is a link between the individual cost of exercising effort and the average effort elicited by the rest of the class. There is a large and growing literature on peer effects (Lavy, Silva, & Weinhardt, 2009; Carrel, Fullerton, & West, 2009; Ding & Lehrer, 2005; Sacerdote, 2001 for the United States). The presence of peer effects offers a reason that it may be efficient to increase continuation tuition in order to modify students’ incentives.

Moreover, by postponing graduation, students could create a negative externality that produces congestion in the classroom, libraries, and so on, thus affecting the learning process of their colleagues. Although our empirical work is mute on these normative issues, each suggests relevant arguments that using the time profile of tuition to change the speed at which a student learns could be optimal.

²⁵ A related issue, also rarely explored, is the choice between a system, such as in undergraduate U.K. courses, in which almost all students finish on time (because it is fairly easy to get a passing grade) and quality is signaled by grades, and the alternative, more common in continental Europe, in which passing grades are harder to get, thus resulting in delayed graduation.

²⁶ The optimal time profile of tuition has been recently analyzed by Gary-Bobo and Trannoy (2004) in a model in which both students and universities face imperfect information on individuals’ ability.

VII. Conclusion

Our evidence suggests that if university tuition were raised for enrollment years beyond normal degree completion, the probability of late graduation would be reduced. This result could be of interest for academic institutions throughout the world that are concerned with the increasing rate at which students delay the completion of a degree.

We exploited data from Bocconi University, where students are assigned to one of twelve tuition levels on the basis of their declared family income, to implement a regression discontinuity design (RDD) that allows us to compare students with similar family income immediately above or below each discontinuity threshold. These two groups of students pay different tuitions but are otherwise identical in terms of observable and unobservable characteristics determining the probability of late graduation. Using this source of identification, we find that 1,000 additional euros of tuition in the last regular year of the program have a negative causal effect on the probability of late graduation as large as 5.2 percentage points, in a context in which the average probability of late graduation is 80%. Since students in the last regular year arguably use their current tuition to predict their future tuition in case of delayed graduation, we interpret this result as an estimate of the causal effect of continuation tuition on the speed of graduation. Such a tuition increase does not induce more students to drop out, and its effect on the speed of completion does not occur at the expense of the quality of the learning process.

We also discussed why it might be optimal to increase continuation tuition with the goal of changing students’ incentives, inducing them to speed up their studies and graduate on time. We have argued that when students are subsidized, peer effects are important, or congestion externalities are relevant, efficiency considerations suggest that continuation tuition should be raised relative to the marginal cost of providing education. More theoretical research and different data would be needed to explore the robustness of these policy conclusions.

REFERENCES

- Angrist, Joshua D., and Guido W. Imbens, “Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity,” *Journal of the American Statistical Association* 90 (1995), 431–442.
- Angrist, Joshua D., Kathryn Graddy, and Guido W. Imbens, “The Interpretation of Instrumental Variables Estimators in Simultaneous Equations Models with an Application to the Demand for Fish,” *Review of Economic Studies* 67 (2000), 499–527.
- Angrist, Joshua D., Daniel Lang, and Philip Oreopoulos, “Incentives and Services for College Achievement: Evidence from a Randomized trial,” *American Economic Journal* 1 (2009), 136–163.
- Angrist, Joshua D., and Victor Lavy, “Using Maimonides Rule to Estimate the Effect of Class Size on Scholastic Achievement,” *Quarterly Journal of Economics* 114 (1999), 533–575.
- , “The Effect of High School Matriculation Awards: Evidence from Randomized Trials,” NBER working paper no. 9389 (2002).
- Booth, Alison L., and Stephen E. Satchell, “The Hazards of Doing a PhD: An Analysis of Completion and Withdrawal Rates of British PhD

Fn 25

Fn 26

COLLEGE COST AND TIME TO COMPLETE A DEGREE

13

- Students in the 1980s," *Journal of the Royal Statistical Society. Series A (Statistics in Society)* 158 (1995), 297–318.
- Bound, John, Michael Lovenheim, and Sarah Turner, "Understanding the Increased Time to the Baccalaureate Degree," mimeograph, University of Michigan (2006).
- Bowen, William G., and Neil L. Rudenstine, *In pursuit of the PhD* (Princeton, NJ: Princeton University Press, 1992).
- Brunello, Giorgio, and Rudolf Winter-Ebmer, "Why do Students Expect to Stay Longer in College? Evidence from Europe," *Economics Letters* 80 (2003), 247–253.
- Carrel, Scott E., Richard L. Fullerton, and James E. West, "Does Your Cohort Matter? Measuring Peer Effects in College Achievement," *Journal of Labor Economics* 27 (2009), 339–464.
- Checchi, Daniele, "Formazione e Percorsi Lavorativi dei laureati dell'Università degli Studi di Milano" Università Statale of Milan, working paper no. 14 (2002).
- Canadian Association for Graduate Studies, "The Completion of Graduate Studies in Canadian Universities: Report and Recommendations" (2003).
- Dearden, Lorraine, Carl Emmerson, Christine Frayne, Costas Meghir, Karl Ashworth, Jay Hardman, Yvette Hartfree, Sue Maguire, Sue Middleton, and Debbi Smith, "Education Maintenance Allowance: The First Two Years. A Quantitative Evaluation," U.K. Department for Education and Skills research report no. 352 (2002).
- Ding, Weili, and Steven Lehrer, "Do Peers Affect Student Achievement in China's Secondary Schools?" mimeograph, Queen's University (2005).
- Dornbusch, Rudi, Andrea Gentilini, and Francesco Giavazzi, "Italian Labor Force Participation: Disguised Unemployment on Campus," unpublished manuscript, Igiar-Bocconi University (2000).
- Dynarski, Susan, "The Behavioral and Distributional Implications of Aid for College," *American Economic Review* 92 (2002), 279–285.
- "Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion," *American Economic Review* 93 (2003), 279–288.
- "Building the Stock of College-Educated Labor," NBER working paper no. 11604 (2005).
- Ehrenberg, Ronald G., and Panagiotis G. Mavros, "Do Doctoral Students' Financial Support Patterns Affect Their Times-to-Degree and Completion Probabilities?" *Journal of Human Resources* 30 (1995), 581–609.
- Filippin, Antonio, and Andrea Ichino, "Gender Wage Gap in Expectations and Realizations," *Labour Economics* 12 (2005), 125–145.
- Garibaldi, Pietro, Francesco Giavazzi, Andrea Ichino, and Enrico Rettore, "College Cost and Time to Complete a Degree: Evidence from Tuition Discontinuities," NBER working paper no. 12863 (2007).
- Gao, Hong, "Examining the Length of Time to Completion at a Community College," paper presented at the Annual Meeting of the Southern Association for Institutional Research, Baton Rouge, LA. (October 2002).
- Gary-Bobo, Robert J., and Alain Trannoy, "Efficient Tuition Fees, Examinations, and Subsidies," CES-Ifo working paper no. 1189 (2004).
- Gneezy, Uri, and Aldo Rustichini, "A Fine Is a Price," *Journal of Legal Studies* 29 (2000), 1–18.
- Groen, Jeffrey, George Jakubson, Ronald G. Ehrenberg, Scott Condie, and Albert Yung-Hsu Liu, "Program Design and Student Outcomes in Graduate Education," NBER working paper no. 12064 (2006).
- Hahn, Jinyong, Petra Todd, and Wilbert van der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica* 69 (2001), 201–209.
- Häkkinen, Iida S., and Roope Uusitalo, "The Effect of a Student Aid Reform on Graduation: A Duration Analysis," Uppsala University, Department of Economics working paper no. 8 (2003).
- Heineck, Martin, Mathias Kifmann, and Normann Lorenz, "A Duration Analysis of the Effects of Tuition Fees for Long Term Students in Germany," University of Konstanz, Centre for European Economic Research discussion paper no. 5 (2006).
- Hoffer, Thomas B., and Vincent Welch, "Time to Degree of U.S. Research Doctorate Recipients," National Science Foundation, Directorate for Social, Behavioral, and Economic Science Info brief no. 312 (2006).
- Kane, Thomas J., "A Quasi-Experimental Estimate of the Impact of Financial Aid on College-Going," NBER working paper no. 9703 (2003).
- Lavy, Victor C., Olmo Silva, and Felix Weinhardt, "The Good, the Bad and the Average: Evidence on the Scale and Nature of Ability Peer Effects in Schools," NBER working paper no. 15600 (2009).
- Lee, David S., "Randomized Experiments from Non-Random Selection in U.S. House Elections," *Journal of Econometrics* 142 (2006), 675–697.
- Leslie, Larry L., and Paul T. Brinkman, "Student Price Response in Higher Education: The Student Demand Studies," *Journal of Higher Education* 58 (1987), 181–204.
- Leuven, Edwin, Hessel Oosterbeek, and Bas van der Klaauw, "The Effect of financial rewards on Students' Achievement: Evidence from a Randomized Experiment," mimeograph, University of Amsterdam (2006).
- McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics* 142 (2008), 698–714.
- OECD, *Education at a Glance 2002* (Paris: OECD, 2002).
- Office of Scientific and Engineering Personnel and Policy and Global Affairs at the National Academies, "On Time to the Doctorate: A Study of the Lengthening Time to Completion for Doctorates in Science and Engineering" (1990).
- Sacerdote, Bruce, "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *Quarterly Journal of Economics* 116 (2001), 681–704.
- Siegfried, John J., and Wendy Stock, "So You Want to Earn a Ph.D. in Economics? How Long Do You Think It Will Take?" *Journal of Human Resources* 36 (2001), 364–378.
- State of Illinois, Illinois Board of Higher Education, "Persistence, Completion, and Time to Degree," mimeograph (1999).
- U.C. Davis, Office of Resource Management and Planning, "Undergraduate Time to Degree Completion Rates by College and Division," issue report (2004).
- U.S. Department of Education, National Center for Education Statistics, "The Condition of Education 2003," NCES 2003-067 (Washington, DC: U.S. Department of Education, 2003).
- van der Klaauw, Wilbert, "Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach," *International Economic Review* 43 (2002), 1249–1287.
- van Ours, Jan C., and Geert Ridder, "Fast Track or Failure: A Study of the Graduation and Dropout Rates of Ph.D. Students in Economics," *Economics of Education Review* 22 (2003), 157–166.