

PHD THESIS DECLARATION

The undersigned

SURNAME *Fano*

FIRST NAME *Shira*

PhD Registration Number *1538354*

Thesis title: *Essays on Labor Economics and Education*

PhD in *Economics and Finance*

Cycle *XXVI*

Candidate's tutor *Professor Tito Boeri*

Year of thesis defence *2016*

DECLARES

Under her responsibility:

- 1) that, according to Italian Republic Presidential Decree no. 445, 28th December 2000, mendacious declarations, falsifying records and the use of false records are punishable under the Italian penal code and related special laws. Should any of the above prove true, all benefits included in this declaration and those of the temporary embargo are automatically forfeited from the beginning;
- 2) that the University has the obligation, according to art. 6, par. 11, Ministerial Decree no. 224, 30th April 1999, to keep a copy of the thesis on deposit at the "Biblioteche Nazionali Centrali" (Italian National Libraries) in Rome and Florence, where consultation will be permitted, unless there is a temporary embargo protecting the rights of external bodies and the industrial/commercial exploitation of the thesis;
- 3) that the Bocconi Library will file the thesis in its "Archivio istituzionale ad accesso

aperto” (institutional registry) which permits online consultation of the complete text (except in cases of a temporary embargo);

4) that, in order to file the thesis at the Bocconi Library, the University requires that the thesis be submitted online by the candidate in unalterable format to Società NORMADEC (acting on behalf of the University), and that NORMADEC will indicate in each footnote the following information:

- thesis *Essays on Labor Economics and Education*;
- by *Fano Shira*;
- defended at Università Commerciale “Luigi Bocconi” – Milano in *2016*;
- the thesis is protected by the regulations governing copyright (Italian law no. 633, 22nd April 1941 and subsequent modifications). The exception is the right of Università Commerciale “Luigi Bocconi” to reproduce the same for research and teaching purposes, quoting the source;

5) that the copy of the thesis submitted online to NORMADEC is identical to the copies handed in/sent to the members of the Thesis Board and to any other paper or digital copy deposited at the University offices, and, as a consequence, the University is absolved from any responsibility regarding errors, inaccuracy or omissions in the contents of the thesis;

6) that the contents and organization of the thesis is an original work carried out by the undersigned and does not in any way compromise the rights of third parties (Italian law, no. 633, 22nd April 1941 and subsequent integrations and modifications), including those regarding security of personal details; therefore the University is in any case absolved from any responsibility whatsoever, civil, administrative or penal, and shall be exempt from any requests or claims from third parties;

7) that the PhD thesis is not the result of work included in the regulations governing industrial property, was not produced as part of projects financed by public or private bodies with restrictions on the diffusion of the results, and is not subject to patent or protection registrations, and therefore not subject to an embargo;

Date *January 30th, 2016*

SURNAME *Fano*

FIRST NAME *Shira*

Contents

Introduction	9
1 The effects of Facebook discussions on academic performance	11
1.1 Introduction	12
1.2 The course and the page	15
1.3 The data	15
1.3.1 Survey data	16
1.3.2 Administrative data	17
1.3.3 External validity	19
1.4 Activity and math grade	20
1.4.1 Basic empirical model	20
1.4.2 Robustness tests	21
1.4.3 Regression results: activity and math grade	22
1.5 Activity and passing probability	25
1.6 Quasi-experimental design	27
1.6.1 Difference-in-difference strategy	27
1.6.2 Difference-in-difference estimates	30
1.7 Concluding remarks	35
Appendices	37
1.A Descriptive statistics	37
1.B Additional material	40
2 Dual labor markets, workers' effort and labor productivity	51
2.1 Introduction	52
2.2 The model	56
2.3 Computational results	59
2.3.1 Payoffs	60

2.3.2	Strategies	61
2.3.3	Effort and productivity	64
2.4	Robustness checks	67
2.5	Concluding remarks	69
Appendices		73
2.A	Empirical evidence	73
2.B	Additional simulation results	74
3	After the crisis, are workers allocated to more productive firms? Evidence from the CompNet dataset	79
3.1	Introduction	80
3.2	The CompNet Database	82
3.3	Labor productivity	84
3.4	Measuring allocative efficiency	85
3.4.1	Time series	87
3.4.2	Sector analysis	88
3.5	Did the crisis affect allocative efficiency?	90
3.6	Allocative efficiency and sector shares	92
3.7	Conclusion	95
Appendices		97
3.A	Sector classification	97
3.B	Country specific notes	98
3.C	Common sector weight computation	99
3.D	Additional material	100
Bibliography		105
Acknowledgements		117

List of Figures

1.A.1	Comparison between observable characteristics in the treated and the control group used in the difference-in-difference.	37
1.B.1	Residual analysis of the OLS model.	40
1.B.2	Comparison between treated and control group observable characteristics.	40
2.1	Share of temporary contract for selected OECD countries in the period 1980–2015 (<i>left</i>) and time series of labor productivity in the years 2000–1014 (<i>right</i>). Data source is OECD database STAN	53
2.1	Time series of average payoffs for different values of the share of permanent contracts, from bottom to top $P = \{0.1, 0.2, \dots, 1\}$	60
2.2	Workers' average strategy at the end of a simulation for all possible values of the share of permanent contracts in $P \in \{0.1, 0.2, \dots, 1\}$	63
2.3	Workers' average effort (<i>left</i>) and probability of exerting minimum/maximum effort (<i>right</i>) for each value of the share of permanent contracts $P \in \{0.1, 0.2, \dots, 1\}$	65
2.4	Share of permanent contracts and labor productivity under two different assumptions on permanent workers' behaviour. The red line is the mean and black is plus/minus 1 standard deviation.	66
2.1	Temporary workers' average strategy for $P = 0.2$ (<i>left</i>) and $P = 0.4$ (<i>right</i>) for three different values of the parameter $\lambda = \{5, 10, 20\}$ determining the speed of the reinforcement learning process.	68
2.2	Share of permanent contracts P and average effort for three values of the effort cost parameter $\alpha = \{0.05, 0.10, 0.15\}$	69
2.A.1	Share of temporary contract (<i>top</i>) share of temporary contracts for the age class 15–24 (<i>bottom</i>); both excluding Spain.	73
2.B.1	Workers' average strategy at the end of a simulation for all possible values of the share of permanent contracts in $P \in \{0.1, 0.2, \dots, 1\}$ when $\alpha = 0.15$, the other parameters are those in the baseline case.	76

2.B.2	Share of permanent contracts and labor productivity when permanent workers supply minimum (<i>left</i>) and 0.5 (<i>right</i>) effort; parameters of the baseline case.	77
3.1	Labor productivity in Belgium (right) and Spain (left) in tradable and non-tradable sectors.	85
3.1	OP-gap in Belgium (right) and Spain (left) in tradable and non-tradable sectors.	86
3.2	Time series of the OP-gap in Belgium and Spain in 1995-2011.	88
3.3	Allocative efficiency by sector, before/after the crisis.	89
3.1	Allocative efficiency and market shares.	95

List of Tables

1.1	Selected survey questions and answers, $N = 236$	17
1.2	Comparison between averages of students in the sample and enrolled from 2009 to 2014 (1141 subjects) in the same Economics degree. The third column provides the p -value of Kolmogorov-Smirnov test.	20
1.3	Effect of activity on math grade, OLS model and robustness tests.	23
1.4	Probit model for the effect of visualizations on passing probability	26
1.5	Comparison between the treated group T and the control group C	27
1.6	Difference-in-difference for the effect of Facebook activity on math grade.	31
1.7	Difference-in-difference for effect of activity on math grade (all T sample)	33
1.8	Difference-in-difference for effect of activity on math grade using administrative data also for T_{12} , so including also students who might have not attended.	34
1.A.1	Comparison between math grade and GPA in the treated and the control subgroups.	38
1.A.2	Comparison between sex and high school final mark in the T and C groups.	38
1.A.3	Comparison between teaching evaluations.	38
1.A.4	Comparison between the 2011 and 2012 sample, administrative data. Column 3 is the p -value of a t -test, where the null hypothesis is $\bar{x}_{2011} = \bar{x}_{2012}$	39
1.A.5	Comparison between the 2011 and the 2012 survey, $217+138=236$ students. Column 3 is the p -value of a t -test, where the null hypothesis is $\bar{x}_{2011} = \bar{x}_{2012}$	39
1.A.6	Summary statistics, $N=236$ observations (102+134).	43
1.B.1	OLS Regression Results	44
1.B.2	Estimates of the selection equation of the Heckman model, 236 observations.	45
1.B.3	Heterogeneous effects of activity on Facebook use	45

1.B.4	Effect of visualizations on passing probability	46
1.B.5	Activity and math grade, alternative activity measures.	47
1.B.6	Activity and passing probability	48
1.B.7	Comparison between the treated group T and the control group C	49
1.B.8	Comparison between the treated group T and the control group C	49
1.B.9	Difference-in-difference estimates, students of previous years as control group.	50
2.1	Summary of reforms implemented between 1980 to 2007 in employment protection legislation (EPL), unemployment benefits (UB), active labor market programs (AP), employment conditional incentives (ECI) and early retirement plans (ER); source Boeri (2011).	52
2.1	Description and value of the parameters used for the simulations.	59
2.1	Average effort for all values of the share of permanent contracts $P = \{0.1, 0.2, \dots, 1\}$ with three values of the learning parameter $\lambda = \{5, 10, 20\}$	68
2.B.1	Average strategy learnt by all workers at the end of a simulation for different values of the share of $P = \{0.1, 0.2, \dots, 1\}$	74
2.B.2	Average productivity for different values of the share of permanent contracts $P = \{0.1, 0.2, \dots, 1\}$, baseline case.	74
3.1	Effect of the crisis on the OP Gap in Belgium and Spain	91
3.2	Effect of the financial crisis on OP Gap, tradable/non-tradable	93
3.3	Financial crisis and OP Gap, robustness tests	94
3.D.1	Effect of the crisis on the OP gap in tradable/non-tradable sectors	101
3.D.2	OP-gap	102
3.D.3	Average Firm Size, Manufacturing sectors	103
3.D.4	Average labor productivity, Manufacturing sectors	104

Introduction

This thesis is a collection of three independent contributions in the fields of labor economics and education, and it focuses on topics at the core of economic debate using methods and tools of econometrics and computational economics.

The first chapter is a co-authored work with Paolo Pellizzari and it is an empirical paper related to education. We study the effects of using a Facebook page exclusively devoted to a first year mathematics course for economics students in a large Italian public university. Posts and discussions supported traditional face-to-face lectures and students could freely post queries and get help from professors and peers. We exploit a quasi-natural experiment to compare the performance of students having access to Facebook with that of a large sample of similar students who were not offered the support page in another branch of the university. Difference-in-difference estimates show that students who could access online discussions gain on average 1 additional point out of 30.

In the second chapter I study the effect of the use of temporary contracts on workers' incentives and, as a consequence, on labor productivity. I implement an agent-based model as it explicitly allows me to take into account the interaction among temporary workers competing for permanent contracts. The main result is that when the share of available permanent contracts is low, workers do not bet on their conversion and supply low effort. As the share increases workers exert higher effort but, when it is too high, they have the incentive to shirk since they are too confident of being confirmed. As a consequence, the relationship between the share of permanent contracts and labor productivity has an inverted-U-shape.

The third chapter is part of a project I was involved in at the Research Division of the European Central Bank, within the CompNet Project (Competitiveness Research Project) aimed at investigating the determinants of competitiveness. I study the effect of the financial crisis on the OP-gap, a measure of allocative efficiency focusing on Belgium and Spain. The main result is that the financial crisis had a positive and significant effect on allocative efficiency only in Spain, not in Belgium and, after the crisis, Spanish sectors were able to further improve the allocation of resources to more productive firms.

In each chapter a conclusion summarizes original results and an appendix collects additional tables and figures not discussed in detail in the text, references are collected at the end of the thesis.

Tesi di dottorato "Essays on Labour Economics and Education"
di FANO SHIRA

discussa presso Università Commerciale Luigi Bocconi-Milano nell'anno 2016

La tesi è tutelata dalla normativa sul diritto d'autore(Legge 22 aprile 1941, n.633 e successive integrazioni e modifiche).

Sono comunque fatti salvi i diritti dell'università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte.

Chapter 1

The effects of Facebook discussions on academic performance

Shira Fano
Paolo Pellizzari

Abstract

In this paper we investigate the effects of using a Facebook page exclusively devoted to a first year Mathematics course in a large Italian public university. Posts and discussions supported traditional face-to-face lectures and students could freely post queries and get help from professors and peers. We use a newly constructed dataset to measure how this influences the grade they achieved and the probability of getting a passing grade.

Firstly, we find that active students, who read and post more often, perform slightly better than non active ones, who mainly read the content, but the effect is not significant. However, other measures of activity, such as the frequency of visualization of the page significantly increase the probability of earning a passing grade, after controlling for students' characteristics and robust *ex-post* measures of ability.

Secondly, we exploited a quasi-natural experiment to compare the performance of students having access to Facebook with that of a large sample of similar students who were not offered the support page in another branch of the university. Difference-in-difference estimates show that students who could access online discussions gain on average 1 additional point out of 30. The effect is, hence, significant but rather small and of possibly limited practical relevance.

Keywords: Facebook, virtual discussions, academic performance, quasi-natural experi-

ment, difference-in-difference.

1.1 Introduction

In the last decade online platforms, digital tools and interactive material are increasingly being used to support traditional teaching methods, Xu and Jaggars (2013). Online learning, in some form, will surely be an increasingly important component of university education, Joyce et al. (2014). This raises the issue of analysing in detail the effect of these new tools on students' satisfaction and academic performance.

Sharing the common idea of using technology and web based applications to support traditional teaching methods, several different tools have been developed and are rapidly spreading in higher education. For example, MOOCs (Massive Open Online courses) are courses entirely taught online, with open access and a potentially unlimited number of participants, see Haber (2014). They were first introduced in 2008 and, although low completion rates are typically an issue, they have become a popular mode of instruction since 2012. Students typically learn by watching recorded lectures, reading material and working on weekly assignments. Blended learning is another form of technology driven tool developed to support education. Although there is no consensus on a unique definition of it, blended learning encompasses several forms of hybrid teaching methods characterized by a combination of both face-to-face learning and online based activities, Bonk and Graham (2012). Flipped learning is a form of blended learning in which students typically use videos prepared by the instructor to study on their own the topics and then discuss and solve exercises in class with colleagues and the instructor, while in traditional lectures the opposite usually occurs.

In this paper, we assess an experiment that consisted in using a Facebook page to support traditional face-to-face lectures of a standard first year Mathematics course for students in Economics at a large Italian public university. Our contribution is to quantify the effects of activity on the page on students' performance, measured in terms of the grade achieved on the math exam and the probability of passing the exam. We do so by controlling for various student characteristics including demographic variables and several measures of ability, such as grade point average and number of credits achieved along the academic career.

Moreover, we take advantage of the fact that Mathematics is a compulsory course in the university where the experiment took place and the same course is taught in a different town to otherwise identical cohorts. We compare the performance of students who could access the Facebook page with the one of similar students of parallel courses who were

not offered a specific support page. In particular, our key methodological contribution is to use a difference-in-difference approach to compare students' math grades in the two samples.

One of the characteristic feature of our experiment is that we choose to use the social network Facebook to host academic discussions, securing a high participation that may not be reached using other tools. For a comprehensive review of the use of Facebook by students and teachers see Hew (2011).

A large body of literature has recently developed trying to quantify the effect of alternative teaching methods on students' performance; but only few studies are conducted in a rigorous manner using quasi-experimental designs or controlling for student characteristics, Means et al. (2009). There is no clear answer on which method is able to produce better learning outcomes, with different studies arriving to opposite conclusions. Two early examples are Navarro and Shoemaker (2000) and Brown and Liedholm (2002); the first presents evidence that students taking an introductory macroeconomics online course have significantly better test scores than students taking the same course in a live lecture format; the second reports the opposite result for students taking an introductory microeconomics class.

There are several papers investigating the effects of video lectures, as opposed to live ones, on students. Figlio et al. (2013) is an insightful work comparing classes where undergraduate students are randomly assigned to attend either live or internet-broadcast lectures. They find no significant difference in the two groups with male and lower-achieving students performing significantly better in live classes. Xu and Jaggars (2013) is another example where the authors estimate the impact of online versus face-to-face lectures on students' performance, using an instrumental variable approach. Their analysis yields robust negative estimates for online learning: attending a particular course in an online format rather than face-to-face would decrease students' likelihood of completing the course and lower their final mark.

Differently from these examples, our Facebook page is more similar to a blended learning experiment as the page is conceived as a tool to support face-to-face lectures, and not to substitute them; in this respect our paper is similar to Bowen et al. (2014), Joyce et al. (2014) and Kwak et al. (2014).

Bowen et al. (2014) examines the performance of students in an introductory statistics class held on six public university campuses. The authors compare the performance of students attending a traditional class with the performance of students whose class material is delivered online supplemented by one weekly class meeting. The study reports no overall difference in performance, measured by their exam grade. Joyce et al. (2014)

is a similar recent study that compares the performance of students who met once per week in a hybrid format, with the one of students in a traditional lecture format of the same class, who met twice per week. The authors show that on average students in the hybrid format scored around 2.5 percentage points less than students in the traditional format, a modest effect according to the authors. Our paper is also related to Kwak et al. (2014) as the authors use a difference-in-difference approach to evaluate the impact of blended learning on students' performance. Their empirical analysis suggests that students' performance is not affected by the course delivery mode when considering only quizzes associated to material covered by blended learning; however when considering an entire course they find a negative effect of blended learning on performance.

As far as the general use of Facebook is concerned, both positive and negative effects on student performance have been documented, with the majority of them tending to be negative. Kirschner and Karpinski (2010), for example, find that an increased activity on Facebook is correlated with significantly lower grade point average and reduced students' time dedicated to studying. The authors claim that multitasking may play a negative role when students switch from academic activities to other online actions. Pasek et al. (2009) criticize the previous results and, using three different datasets, find no significantly negative effect of Facebook use on students' grades, controlling for student characteristics such as ethnicity and socio-economic status; on the contrary in one of the three dataset they find a weak positive effect.

The previous two papers study the effect of Facebook on academic performance, when it is used for general purposes of leisure and networking, but not specifically related to academic activity. The key contribution of our paper is instead to investigate the effect of activity on the social network, when it is specifically used as a support to traditional face-to-face lectures. Two qualitative examples of the effect of Facebook, when it is used as a supporting teaching tool, are Bosch (2009) and Kabilan et al. (2010). The first, presents results of an experimental use of Facebook at the University of Cape Town, and it draws positive conclusions on the use of Facebook as an additional learning channel. The second, uses Facebook as a supporting tool for an English language course, and it also draws positive conclusions. Over 60% of surveyed students agree or strongly agree on all positive evaluations of Facebook regarding improvements in language writing and reading skills, motivations and confidence.

The rest of the paper is organized as follows. Section 1.2 describes the course and the page and section 1.3 the data. Sections 1.4, 1.5 and 1.6 present methods and results for the econometric models used; section 1.7 concludes.

1.2 The course and the page

The course “Mathematics” is a compulsory first-year course for undergraduate students enrolled in one of the Economics degrees offered by the university. Instructors are tenured faculty members who typically have taught the course for many years. Topics covered included calculus, functions of one and two variables, some linear algebra and elementary notions of financial discounting. Lectures are held 3 times per week and last 90 minutes, for a total of 30 meetings. In the two years examined in this paper, classes were crowded and more than 150 students attended the lectures. Students could choose to take the exam in two ways: they could either sit two midterms in November and December splitting in half the material covered, or one unique final held in two occasions in January. The course took place in the first semester, between mid September and mid December.

At the beginning of the course a Facebook page was activated and administered by the instructors, with a twofold aim. First, it offered a quick and effective way to provide organizational and administrative information on practical aspects of the course such as the timetable of lectures and availability of office hours. Secondly, the page provided an additional tool for students who could ask both their colleagues and the instructor questions on difficult topics, additional explanations on the teaching material and help to solve exercises. In fact, mathematics is perceived as a difficult subject by Economics students, coming from different backgrounds and with different quantitative skills. Moreover, the page allowed shy students, who would not dare to ask questions in a crowded class, to post their doubts and get quick feedback. On a weekly basis, the syllabus was updated and posted on the page, together with additional exercises and mock exams allowing students to check their proficiency on required basic tasks.

1.3 The data

We use a newly constructed dataset by merging information from two different sources: a survey filled out by students at the end of the course, before any exam, and administrative data retrieved from the University’s official records.

The Facebook page was activated for the first time in the fall semester of 2011 and the same experiment was repeated the following year. Both years, an online administrated survey was filled out by students on a voluntary basis. The two courses were identical in terms of syllabus, provided material and instructors; we compared survey answers and students’ observable characteristics in the two waves and found no meaningful differ-

ence.¹ Therefore, we pool data from the two waves and overall we have a sample of 355 undergraduate students, 217 and 138 from 2011 and 2012 respectively.

In the 2011 survey providing personal information (name or student ID) was voluntary, while this information was required the following year, under the standard agreement that only aggregate data would have been disclosed. The knowledge of the identity of the respondent is necessary to match the answers from the survey with administrative data and retrieve students' characteristics and academic records. In the econometric analysis we focus on the restricted sample of $N = 236$ students who provided identification and were successfully matched with the administrative data (102 and 134 are drawn from the samples of 2011 and 2012, respectively). The overall sample is homogeneous with respect to the subject studied (Economics), status (undergraduate) and ethnicity (Caucasian, with a handful of exceptions).

1.3.1 Survey data

Information collected with the survey includes sex, level of activity on the online platform, perceived usefulness of the page and information on the amount of time spent on the Facebook page and studying for the exam. The survey was screened by a statistician to eliminate possible biases in the way the questions were worded. Table 1.1 reports some of the most relevant questions asked in the survey and the frequency of chosen options.

Two thirds of the students defined themselves as “not very active” and approximately one third as “fairly active”, only 3 and 4 students picked “one of the most active” and “one of the least active”, respectively. In the econometric analysis we merged them with the other two factors “fairly active” and “not very active”. Visualizations were frequent: 26 percent of students visualized the page several times per day, 47 percent once per day and 23 percent few times a week. Students highly appreciated the experiment as 51 percent valued the usefulness of the page at the top level and 25 and 16.9 percent of students choose 9 and 8 respectively. Students' expected gains from the page to be significant and more than two thirds of students thought they would gain 2 points (33.9 percent) or more (38.1 percent).²

In the 2012 version of the survey, to better understand students' time allocation, four question were added: “On average, how much time do you spend daily on Facebook?” (median=45 minutes, range=5-200), “Which fraction of time spent on Facebook is ded-

¹The only statistically significant difference is that in 2012 visualizations of the Facebook page were more frequent, with less students visualizing the page few times a week and more several times per day.

²In Italy, the grade is routinely expressed in thirtieths: failing to reach 18/30 forces the student to retake the exam. We further assume that 32 is equivalent to the top grade 30/30 *cum laude*.

icated to studying?” (zero=3%, 0-10%=31%, 10-20%=30%, 20-50%=19%, more than 50%=17%), “On average, how much time did you spend daily on the [academic] page?” (median=10 minutes, range=1-100) and “On average, how much time did you daily spend studying for the math exam?” (median=120 minutes, range=2-480).

Table 1.1: Selected survey questions and answers, $N = 236$.

Question	Options	Percent
Which kind of user are you?	fairly active (read& post few times)	30.9
	not very active (mainly read)	66.1
	one of the least active	1.7
	one of the most active	1.3
How often did you visualize the page?	when notifications appeared	0.8
	before the exam	3.0
	few times a week	23.3
	once a day	46.6
Do you think the page was usefull?(1-10)	several times a day	26.3
	useful 1	0.4
	useful 6	0.8
	useful 7	5.9
	useful 8	16.9
Evaluate the usefulness in terms of final grade out of 30: which is the added value?	useful 9	25.0
	useful 10	51.0
	no difference	3.8
	1 point	15.3
	2 points	33.9
	more than 2 points	38.1
	don't know	8.9

1.3.2 Administrative data

The second source of information is data retrieved from the University’s official records, after obtaining written permission from ADiSS (Area Didattica e Servizi agli Studenti), the Teaching and Students’ Services Division, which allowed the retrieval of selected infor-

mation and prescribed strict anonymization of personal information. This data includes demographic information such as sex and province of residence; information on the students' educational experience prior to entering university such as type of high school attended and final mark; and university records such as track ("corso di laurea") chosen, year of enrolment, exams passed, grades and credits achieved. This data is both matched with the survey data and used as a source to build the control group in the quasi-experimental approach.

We consider several indicators of students' performance, used in the econometric analysis to control for student "quality" and proxy unobserved ability. Since mathematics is a first year course, at the time of the experiment on-track students had passed few exams and collected few credits, preventing us from having adequate and robust measures of their ability. Therefore, we waited and tracked their performance up to July 2014 and computed their grade point average and number of credits, obtaining robust *ex-post* measures of student ability.

Grade point average is a synthetic measure of students' performance; in our sample the median is 23.8 (mean 24.1) and 19 and 29 are the minimum and maximum.

The number of credits achieved per year, compared to grade point average, gives additional information on students' speed, see Brugiavini et al. (2014).³ In the Italian system students are not required to sit the exams at the end of the course, and some students happen to procrastinate taking exams. In the mathematics course there were both on-track students, enrolled in the first year, and out-of-track students, enrolled in the previous years, who decided to attend the course after some time. To take this into account, we standardize the total number of credits gained and obtain a measure that is comparable across different enrolment years.⁴ The median number of standardized credits, ranging from 0 to 1 by construction, is 0.54. The correlation between this measure and grade point average is 0.5. This means that there are "fast" students with relatively low GPA and "slow" students with relatively high GPA.

We consider two additional performance measures: the high school grade and the score obtained in a formal mathematics entry test that first-year students take upon enrolment. The high school grade is achieved at the end of a 5-year schooling period required by law to

³In the Italian system each exam is worth a certain amount of credits, depending on the work load required, and students need to earn 180 credits to obtain the undergraduate degree.

⁴For every year of enrolment we rank students according to the total number of credits they achieved and compute the percentiles of each distribution; we assign to each student the corresponding percentile. For each year, the student with the least (most) credits will get 0 (1). A value of, for example 0.3, means the student performed better than 30 percent of the students enrolled in the same year, in terms of number of credits achieved.

access university. It is a general and noisy measure of scholastic skills, as it is a summary of the performance in several distinct disciplines. In our sample, the median high school grade is 79.65, on a scale from 60 to 100. Regarding the entry test, failing to achieve a given threshold temporarily halts the admission process, until the student gains additional credits by attending special-purpose refresher courses in mathematics and proves his/her proficiency. In our sample the median is 20, on a range from 7 to 36. Compared to the high school grade, the entry test score is a more specific measure of the quantitative and mathematical abilities of the student. However, due to administrative reasons, data on the entry test is only available for the 2011 wave.

The grade achieved in the mathematics exam is the dependent variable used in the econometric analysis. The median grade is 24 and 18 and 30 *cum laude* are the minimum and maximum. We consider relevant in our study only the grades of students who successfully passed the exam in the midterms or in the two examinations held in January, after the students followed the course. We make the assumption that, if the access to online discussions had an effect, it is only for students who actively followed the course and took the exam right after. We believe it is extremely unlikely that students who took the exam in June or later could still benefit of conversations or material posted months heretofore. In the sample 176 out of 236 students (75 percent) succeeded in passing the exam within the January sessions.

1.3.3 External validity

Students decided to fill in the survey on a voluntary basis. This may cause a self-selection problem and therefore our sample may not be representative of the broader population of potential students. If those who appreciated the Facebook page or have strong quantitative skills are more willing to respond to the survey, and these students respond differently to online activity on the Facebook page, then our results cannot be generalized to typical students enrolled in a large introductory math course. The ideal situation would have been to have all students attending the course fill in the survey (and provide their student ID), but this could clearly not be enforced.

As magisterially exemplified in Figlio et al. (2013), we did our best to check whether the students who volunteered to fill in the survey and provided ID are different in observable ways from the entire population of students enrolled in the same Economics degree since 2009. We compare students' performance in terms of: grade point average, number of standardized credits achieved and high school final mark. Table 1.2 summarizes the comparison between the two samples and column 3 is the p -value of a Kolmogorov-

Smirnov test, where the null hypothesis is that the two samples were drawn from the same distribution.

Table 1.2: Comparison between averages of students in the sample and enrolled from 2009 to 2014 (1141 subjects) in the same Economics degree. The third column provides the p -value of Kolmogorov-Smirnov test.

	\bar{x}_{sample}	\bar{x}_{pop}	p -value
GPA	24.08	24.22	0.94
credits std	0.54	0.56	0.22
high school	79.65	77.8	0.14

As can be seen from the table, for the three variables grade point average, standardized number of credits and high school final mark, we cannot reject the null hypothesis. The results of a simpler t -test on the difference of the means yields similar results: p -values for GPA, standardized credits and high school final mark are 0.38, 0.42 and 0.03, respectively. Moreover, to rule out the hypothesis of time trends, we repeated the comparison between distributions using subsamples of the population, made of students enrolled since 2010, 2011 and so on; results qualitatively and numerically do not change.

In summary, the sample appears to be reasonably representative of the overall population of the students or, at least, we see no obvious biases although it was not randomly selected.

1.4 Activity and math grade

1.4.1 Basic empirical model

Our goal is to estimate the effect of students' engagement on the Facebook page on the grade they achieved on the mathematics exam. To explore this issue quantitatively, we start with a linear regression model estimated by ordinary least squares (OLS). We fit a simple linear regression model of the form:

$$grade_i = \beta_0 + \beta_1 \cdot activity_i + \beta_2 \cdot ability_i + X_i' \beta_3 + \varepsilon_i \quad (1.1)$$

where the $grade$ obtained in the mathematics exam is the dependent variable and $activity$ is the main explanatory variable. Therefore, the coefficient β_1 quantifies the effect of students' online activity on the Facebook page on their math grades. X_i' is a row vector of student characteristics and additional controls, including demographic variables and

perceived usefulness of the Facebook page. In addition, we include year fixed effects since the experiment has been repeated for two years.

The most active students could earn better grades thanks to their activity on the Facebook page, which is the possible causal effect we are interested in. It could instead be that better students *ex-ante* are more likely to participate to online discussions *ex-post* thinking this would improve their performance. To take this causality issue into account, we use several measures to control for student ability such as: grade point average, standardized number of credits achieved, score obtained in the university entry test and high school final mark.

The concept of activity is multidimensional and nuanced and, therefore, its measurement is far from trivial. In the baseline case, activity is a dummy variable equal to 1 if the student defined himself as “fairly active (I read and post few times)” or “one of the most active” and it is equal to 0 if the student is “not very active (I mainly read the content)” or “one of the least active”.

Defined in this way, one could interpret the effect of activity as the marginal benefit obtained by students who not only read the posts, but also actively wrote on the Facebook page. Nevertheless, also the students who simply read the content, perhaps the less performing ones, could benefit by spending time on the page or frequently visualizing the information posted by their colleagues. Therefore, as additional measures of activity, we consider the frequency of visualizations and the time spent by students on the Facebook page.

1.4.2 Robustness tests

As a robustness test, we use a Tobit model to account for the fact that the dependent variable, namely the grade achieved in the mathematics exam, is limited both from below and from above by 18 and 30 *cum laude*. If there is a significant number of predicted values below and above the truncation points this may bias the results and in this case the truncated regression is a valid choice.

The OLS and Tobit specifications cannot account for another potential problem. In the previous regressions only the students who successfully achieved a passing grade, greater than or equal to 18, are included in the sample, therefore, using this non-random sample might bias our estimates. To correct for this bias, we use the Heckman selection model, see Heckman (1979), originally used to estimate the determinants of wage offers in the US labor market. In the first stage we estimate the probability of achieving a passing grade using a Probit model. As a predictor for the probability of passing the math exam,

we use the variable standardized number of credits, and expect a positive and significant correlation. The estimates from the selection equation are used to predict the probability of achieving a passing grade for each individual. In the second stage we correct for the selection bias including a transformation of the predicted individual probabilities as an additional explanatory variable in a linear regression with grades as a dependent variable.

1.4.3 Regression results: activity and math grade

The estimates of the OLS model are shown in Table 1.3, together with other models used as robustness checks. The math grade is the dependent variable and activity is the main explanatory variable; it is a dummy equal to 1 if the student is “fairly active” or “one of the most active”, 0 otherwise. We test the relevance of different subsets of regressors and subsequently eliminate the ones that are not statistically significant. The first column shows the estimates of the selected OLS model.

As can be seen from the table, we find weak evidence in favour of the hypothesis that active participation to discussions (i. e., reading *and writing*) helps students to improve their grade at the exam. Regression coefficients indicate that activity has the expected sign: the ones who read and post perform better than not active ones, who mainly read the content, and earn 0.28 additional points, but the magnitude of the difference is small and neither statistically nor practically significant.

As largely expected, students with higher grade point average perform significantly better in the exam, earning 1.2 points for every additional GPA point.

The quantitative and mathematical abilities that students have when they enter university, measured by the grade on the entry test, are positively correlated with the math grade, and the estimated correlation is 0.2. Since the average grade on the entry test is 19.5, the average increase in math grades is approximately 4 points. The dummy for students who did not take the entry test has a similar magnitude. The correlation between *activity* and *GPA* is -0.04 and the correlation between *activity* and students' score in the entry test, *entryz*, is 0.02. Therefore, we believe there is no self-selection into activity driven by ability.

The dummy for students who valued 10 the usefulness of the page is positive and statistically significant: students who ranked the usefulness at the top earned on average 1 additional point, after controlling for the other covariates. Interestingly, students who found the page most useful actually benefited from it. This suggests that not all students are alike and not everyone profits from online discussions in the same way. On the Facebook page short posts and status updates are the main means of communication,

information flows rapidly and, therefore, posts can rapidly sink and are no longer visible in a very short amount of time.

Table 1.3: Effect of activity on math grade, OLS model and robustness tests.

	<i>Dependent variable: math grade</i>			
	<i>OLS</i>	<i>robust linear</i>	<i>Tobit</i>	<i>Heckman selection</i>
activity	0.282 (0.460)	0.381 (0.494)	0.417 (0.513)	0.404 (0.462)
GPA	1.213*** (0.105)	1.261*** (0.113)	1.356*** (0.121)	1.280*** (0.117)
entryz	0.218*** (0.060)	0.237*** (0.064)	0.267*** (0.068)	0.219*** (0.059)
entryno	4.045*** (1.363)	4.747*** (1.464)	4.977*** (1.563)	4.145*** (1.331)
useful10	0.990** (0.433)	1.133** (0.465)	1.125** (0.485)	0.965** (0.423)
Constant	-10.777*** (2.731)	-12.674*** (2.933)	-15.456*** (3.199)	-12.900*** (3.226)
Time FEs	YES	YES	YES	YES
Observations	176	176	176	236
R ²	0.513			0.516
Adjusted R ²	0.495			0.496
F Statistic	29.633***			
Wald Test			166.087***	

Note:

*p<0.1; **p<0.05; ***p<0.01

It might be that students who use smartphones or tablets, or are often online, are

able to ride the information flow and take advantage of it. Instead, students who prefer traditional tools as textbooks, paper and pen may not be able to benefit from this hectic environment and could be distracted from the option to be active online.

We include time fixed effects to take into account that in the sample we have students from the two waves 2011 and 2012 and, as expected, being in one or the other wave has no significant effect on math grades, controlling for the other covariates. The coefficient of determination, in the basic or adjusted form, tells us that the model explains approximately 50% of the variation in the data, which is a reasonable amount for a cross section. The elimination of the non significant variables increased the adjusted- R^2 and did not change the magnitude of the other significant variables. The variables we dropped from the model are: sex, dummies for perceived usefulness of the page equal to 8 and 9, number of standardized credits achieved, expected gains and high school final mark. Moreover, the alternative activity measures, frequency of visualizations and time spent on the page (for the 2012 sample) are not significant.

As preciously explained, we perform several robustness test: the estimated coefficients of a robust regression, a Tobit model and the Heckman selection model are displayed in columns 2–4 of Table 1.3. First, to control for heteroskedasticity, we fit a robust regression model and the estimated coefficients closely resemble those in the first column: the coefficient for activity increases to 0.38 but it is still not statistically significant. Second, we use a Tobit model to take into account that the dependent variable is bounded between 18 and 30 *cum laude*. The coefficients are quite similar to the ones obtained with OLS and the same can be said for the marginal effects (0.40, 1.31, 0.25, 4.8 and 1.08 for the independent variables). Third, we fit a Heckman selection model to take into account the possible selection bias due to the inclusion in the model only of student who achieved a passing grade. In the selection equation, the standardized number of credits achieved is a significant predictor of the probability of passing the math exam. The estimated coefficients of the outcome equation are shown in column 4 of Table 1.3; they closely resemble those of the other specifications suggesting that using only students who successfully achieve a passing grade does not bias the estimates. We repeated the estimation of the Heckman model with different specifications in the first stage using, for example, the grade achieved in the statistics exam as a predictor of the passing probability, since it is another compulsory quantitative course; the results both qualitatively and numerically do not change.

To sum up, our results are quite robust with respect to several specifications, the estimated parameters are similar across different models and the relations among variables look quite stable.

1.5 Activity and passing probability

We perform a second set of regressions to estimate the effect of students' engagement on the Facebook page on the probability of achieving a passing grade on the math exam. We use a Probit model of the form:

$$P(\text{pass}_i = 1 | \text{activity}_i, \text{ability}_i, X_i') = \Phi(\beta_0 + \beta_1 \cdot \text{activity}_i + \beta_2 \cdot \text{ability}_i + X_i' \beta_3) \quad (1.2)$$

where Φ is the c.d.f. of the standard normal and the dependent variable is a dummy equal to 1 if the student passed the math exam, earning 18 or more, and 0 if he failed or did not sit the exam within the two examinations held in January. As in the OLS case, we use different measures of activity, and control for student ability using several *ex-ante* and *ex-post* performance measures. Moreover, we test the robustness of our results by using alternative models for binary dependent variables.

Compared to the OLS regressions, standard errors are smaller and the estimates are more precise, reflecting the bigger sample size. As in the OLS model our baseline measure of activity, the dummy equal to 1 if the student is "fairly active" or "one of the most active" is not significant. Instead, when we use the frequency of visualizations as a measure of activity, we find interesting results: Table 1.4 shows the estimates of the effect of the frequency of visualization of the Facebook page.

In column (1) we consider a model where the independent variables are restricted to the dummies for the number of visualizations (a few times per week, once a day and several times a day) and the grade point average to proxy for unobservable ability.

We find that the frequency of visualizations is positively and significantly correlated with the probability of passing the math exam. Moreover, visualizing the page more often monotonically increases the probability of success, compared to students who visualized the page rarely, i.e., only before the exams. Computing marginal effects, we find that visualizing the page a few times a week, once a day and several times per day increases the probability of passing the exam respectively by 23, 26 and 32 percent.

As expected, grade point average is positively and significantly correlated with the probability of success, as every additional point in students' GPA increases the chances of passing the math exam by 6 percent. In our sample the least and most performing students have a grade point average of 19 and 29 respectively, so the best students have 60 percent higher chances of passing the exam, other things being equal, compared to the least prepared ones.

In the second column of Table 1.4 we included several control variables, as well as time fixed effects to take into account that in the sample we have students of the two different

waves. Results are robust, as the estimated coefficients are similar to those in column (1), both in terms of magnitude and significance. Visualizing the page a few times a week, once a day and several times a day increases the probability of passing the math exam by 22, 24 and 28 percent and the marginal effect of grade point average is still a sizeable 4 percent.

Table 1.4: Probit model for the effect of visualizations on passing probability

<i>Dependent variable: pass (1 if grade \geq 18)</i>		
	(1)	(2)
GPA	0.253*** (0.050)	0.151*** (0.058)
vfewtweek	0.919* (0.489)	0.913* (0.522)
vonced	1.041** (0.471)	0.993** (0.502)
vseved	1.265*** (0.490)	1.141** (0.527)
Constant	-6.357*** (1.252)	-5.259*** (1.405)
Additional controls	NO	YES
Time FEs	NO	YES
Observations	236	236
Log Likelihood	-114.775	-101.978
Akaike Inf. Crit.	239.551	225.956
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

As a further robustness test, we used alternative models for binary dependent variables, such as the logit and the linear probability model. Results are similar to those obtained with the Probit model and the estimates confirm the monotonically increasing relation between the frequency of visualizations and passing probability, both with and without additional covariates.

1.6 Quasi-experimental design

Several parallel math courses are taught at the same university and we exploit this fact to implement a quasi-experimental design to compare the performance of students who had access to online discussions with the one of similar students who did not have access to any support page. In particular, we use a difference-in-difference approach to measure the difference in math grades between the two groups.

1.6.1 Difference-in-difference strategy

We call T the treated group, made of students that took the math exam in the town where the online support page was offered and C the control group made of students in parallel courses who did not have the page. We estimate the effect of active participation to online discussions on the grade achieved in the math exam by comparing the average change over time in math grades for students in the T group with the average change over time for students in the C group. We combine data from the survey and administrative data reaching a sample size of $N = 2380$ students.

The difference-in-difference requires data measured at two or more points in time; we therefore assume that the experiment took place in the 1st and 2nd examinations held in January, as previously discussed, and use the 3rd and 4th examinations as the post experimental period.

The math courses followed by the treated group are by all means similar to the parallel courses attended by the control group. All courses have the same program, textbooks, number of face-to-face lectures, exam mode and availability of office hours. In Table 1.5 we compare students' observable characteristics in terms of math grade, grade point average, number of standardized credits achieved and high school final mark; column (3) is the p -value of a Kolmogorov-Smirnov test.

Table 1.5: Comparison between the treated group T and the control group C .

	\bar{x}_T	\bar{x}_C	p-value
math grade	22.310	22.266	0.937
GPA	23.786	23.925	0.350
credits std	0.528	0.563	0.007
high school	77.847	77.269	0.306

Overall the treated and the control group are similar in observable characteristics: the only statistically significant difference is in the earned standardized credits, but the

magnitude of the discrepancy is small and, if anything, it shows that the students in the control group are slightly better *a priori*, thus reinforcing our findings.⁵ We further investigated the comparability of the samples in two ways.

First, we compared the “quality” of high schools in the two towns where the math courses were held for the treated and the control group. We used an online portal that allows to compare different types of high schools, within a given region, in terms of the performance achieved *ex-post* by high school graduates when they go to university.⁶ Grade point average, credits and an index combining the two are used to compare high schools, and weights are used to allow the comparability across different faculties. This tool should advice students on which are the high schools that better prepare students for university. Results of the simulations done do not show any major concern and, if anything, they suggest that high schools in the town where the math course was activated with the Facebook page are slightly better.

Second, we compared teaching evaluations filled in by students at the end of all courses, both in the track where the experiment took place and in the other economics degrees; the questions asked are identical. We focus on teaching evaluations of the math courses and on an average measure of all courses in the two academic years studied 2011/12 and 2012/13. On a scale ranging from 1 to 5, average satisfaction for the math course taken by treated students is 3.57 and 3.67 respectively in the 2011 and 2012 wave; in the control group it is respectively 3.31 and 3.33. The average evaluation of all courses in the treated group is 3.15 and 3.22 respectively in 2011 and 2012 and 3.18 and 3.25 for tracks in the control group.

The higher appreciation of the math course in the treated group may depend on the quality of instructors but also on the availability of the Facebook page.

In the econometric analysis we control for student ability using the different performance measures, and include time and track (“corso di laurea”) fixed effects as students in the different tracks in the control group may have different quantitative skills or differ in unobservable ways.

The math courses of the treated and the control group have different instructors that may have (slightly) different grading methods. However, the comparison of the math grades of the 1st and 2nd examinations with the ones of the 3rd and 4th examination allows us to net out this potential bias. To identify the causal effect of Facebook activity

⁵Results of a *t*-test comparing means yields similar results; p-values for math grade, GPA, standardized credits and high school mark are 0.827, 0.183, 0.059 and 0.311 respectively.

⁶The website used, developed by the *Fondazione Giovanni Agnelli* is:
<http://www.eduscopio.it/percorso-studenti-scelta-scuola-superiore#mapSlide>

we are implicitly taking the *parallel trend assumption*. In our case, this means that if there is a difference in grades between the January session and the following ones, this difference is the same for all students, both in the treated and the control group, and that each instructor uses the same grading method in the different sessions. We believe this is a reasonable assumption. If self-selection into the course with the Facebook page and students' ability were correlated in an unobservable way, controlling for student characteristics would not be enough to eliminate the selection bias. However, with our quasi-experimental design we can eliminate this bias as, at the time of enrolment, students did not know that just in the T group the experiment would take place and, in any case, students would not chose in which course to enrol depending on this. After controlling for student characteristics and instructor grading methods, and including fixed effects to control for the different waves and tracks, it can be argued that the remaining difference between the T and the C group is mainly due to the availability of a discussion page and, hence, the difference-in-difference provides a consistent estimate of the (causal) effect of student activity on the Facebook page on their performance.

Figure 1 illustrates the difference-in-difference estimate, representing the effect of on-line discussions.

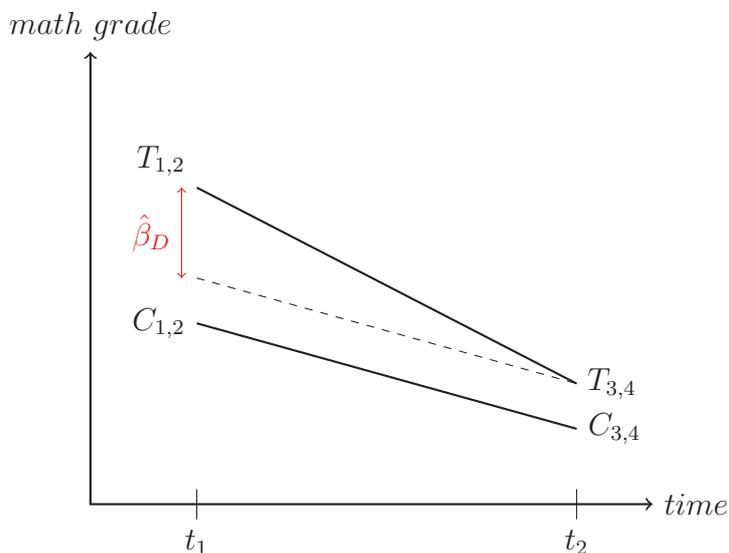


Figure 1: Difference-in-difference estimate.

The treatment (control) group is depicted by the T (C) line. $T_{1,2}$ and $C_{1,2}$ are the average math grades in the treated and control group at t_1 , when the experiment took place, while $T_{3,4}$ and $C_{3,4}$ are average math grades at t_2 . The difference between average math grades at t_1 cannot be imputed entirely to Facebook, say because instructors of the T group may tend to give higher marks or due to other unobservables. The difference

between $T_{3,4}$ and $C_{3,4}$ allows us to net out this potential bias and therefore the difference-in-difference $\hat{\beta}_D$ is a consistent estimate of the (causal) effect of Facebook activity on math grades for the treated group.

We estimate the following difference-in-difference equation:

$$grade_i = \beta_0 + \beta_1 \cdot January_i + \beta_2 \cdot T_i + \beta_D(January_i * T_i) + X_i' \beta_3 + \varepsilon_i \quad (1.3)$$

where $January_i$ is a dummy equal to 1 if student i passed the exam in the 1st or 2nd examination held in January and 0 if he/she passed the exam in the 3rd or 4th examinations in the following session; T_i is a dummy equal to 1 for students in the treated group and 0 in the control group and X_i' is a row vector of additional covariates. The composite variable $January_i * T_i$ is then a dummy variable indicating when $T_i = January_i = 1$, referring to students in the treated group who took the exam in January and potentially benefited from activity on the Facebook page.

1.6.2 Difference-in-difference estimates

Table 1.6 reports the coefficients of several specifications of the difference-in-difference model.

In column (1) we only include the dummies $January$, T and their interaction; due to omitted variable bias the model is not correctly estimated, and the R^2 is very low.

In the model in column (2) we add grade point average as an additional explanatory variable; this greatly improves the explanatory power of the model and changes both the magnitude and the significance of the estimated coefficients. The interaction term $January * T$ is the difference-in-difference estimate that, as previously discussed, measures the effect of activity on the Facebook page on students' math grades for the treated group. The estimated coefficient is positive and statistically significant at 10% and students who are active on the Facebook page and pass the math exam in January gain on average 0.81 supplementary points. Observe that the additional performance of the treated students, 0.322, is statistically and practically insignificant, whereas being treated and taking the exam soon has a much larger marginal effect that can be attributed to the role of online activity.

As expected, grade point average is still a significant variable in explaining the variability in math grades and every additional point in students' GPA increases their grades by approximately 1 point. Students who pass the exam in January earn on average 0.54 points more than students who decide to delay the exam or do not pass it at the first

attempt and are forced to retake it, and the difference is statistically significant.⁷

This is in line with previous findings in the literature, see for example Cappellari et al. (2012), who show that the longer students of Economics wait to take the mathematics exams, the less likely they are to obtain high grades.

Table 1.6: Difference-in-difference for the effect of Facebook activity on math grade.

<i>Dependent variable: math grade</i>				
	(1)	(2)	(3)	(4)
January	1.546*** (0.188)	0.539*** (0.159)	0.411** (0.163)	0.284* (0.162)
T	0.902*** (0.263)	0.322 (0.219)	0.032 (0.266)	-1.941** (0.893)
January*T	0.623 (0.555)	0.814* (0.461)	1.273*** (0.482)	1.275*** (0.481)
GPA		1.058*** (0.032)	1.058*** (0.040)	1.107*** (0.040)
Constant	21.075*** (0.165)	-3.452*** (0.763)	-1.706 (3.077)	-2.976 (3.029)
Additional controls	NO	NO	YES	YES
Course FEs	NO	NO	NO	YES
Time FEs	NO	NO	NO	YES
Observations	2,380	2,380	2,349	2,349
R ²	0.034	0.334	0.353	0.383
Adjusted R ²	0.033	0.333	0.346	0.370
F Statistic	28.275*** (df = 3; 2376)	297.657*** (df = 4; 2375)	52.783*** (df = 24; 2324)	28.523*** (df = 50; 2298)

Additional controls: sex, high school type, high school final mark, standardized credits, dummies for year of enrolment.

Note:

*p<0.1; **p<0.05; ***p<0.01

The model in column (3) checks the robustness of our findings with the inclusion of

⁷We do not record failed attempts to take the exam. If a student did not pass in January, and passed in June, his grade is simply recorded in June.

several additional controls: sex, type of high school attended and final mark, number of standardized credits achieved and dummies for students' year of enrolment. Results obtained are similar to those in column (2), the magnitude of the interaction term increases to 1.2 and significance also increases. The only significant additional covariate is the dummy for "licei" that are high schools typically chosen by students who plan to go to university, and these students obtain on average 0.65 points more.

In the fourth column we add course fixed effects, to take into account the different tracks chosen by students in the control group, and time fixed effects since we consider students of the two waves 2011 and 2012. Results robustly confirm the positive effect of online activity.⁸ The model explains 37% of the variability in the data, and adding fixed effect increases the models' explanatory power by approximately 2%, compared to the specification in column (3).

In the previous regressions we include as treated students who passed the math exam within the January session (sub-sample $T_{1,2}$) only students who defined themselves as "fairly active" or "one of the most active". We perform a second set of regressions including also students who defined themselves as "not very active" or "one of the least active", i.e, all students who filled in the survey, see Table 1.7.

As shown in column (4) the interaction term $January_i * T_i$ is still significant but, as expected, the magnitude drops to 0.94.

This interestingly points out that also students that did not defined themselves as active, but perhaps content themselves to visualize the page, earned higher grades. Another possible interpretation is that the Facebook page generated positive spillovers and also students who did not directly use the page indirectly benefited from it, for example, because they studied with their colleagues that were active on the page.

In the previous set of regressions, we included in the treated group all students that filled in the survey, regardless on whether they defined themselves active or not. Nevertheless, potentially also students who decided not to fill in the survey could have benefited, directly or indirectly, from discussions on the Facebook page. As a final and sharp robustness test, we perform a set of regressions in which we include in the sub-sample $T_{1,2}$ of treated students all students who passed the exam in the January session of the two waves studied, 2011 and 2012, retrieving this data from the administrative records; estimates are reported in Table 1.8.

The interaction term confirms the positive effect of online discussions on performance

⁸The estimated coefficient for the effect of being in the T group is -1.941 but this is because the variable interacts with the course fixed effect that identifies the same subgroup, and its estimated coefficient is 1.898. Therefore the marginal effect of being in the T group is basically 0, as in column (3).

but, as expected, the magnitude of the effect is smaller than in the previous specifications.

Table 1.7: Difference-in-difference for effect of activity on math grade (all T sample)

	<i>Dependent variable: math grade</i>			
	(1)	(2)	(3)	(4)
January	1.546*** (0.189)	0.528*** (0.159)	0.413** (0.162)	0.286* (0.161)
T	0.902*** (0.264)	0.316 (0.219)	0.045 (0.265)	-1.279 (0.837)
January*T	0.261 (0.389)	0.457 (0.321)	0.939*** (0.352)	0.943*** (0.350)
GPA		1.069*** (0.031)	1.078*** (0.039)	1.123*** (0.039)
Constant	21.075*** (0.166)	-3.707*** (0.735)	-2.094 (3.061)	-3.197 (3.017)
Additional controls	NO	NO	YES	YES
Course FEs	NO	NO	NO	YES
Time FEs	NO	NO	NO	YES
Observations	2,501	2,501	2,470	2,470
R ²	0.038	0.346	0.365	0.393
Adjusted R ²	0.037	0.345	0.359	0.380
F Statistic	33.090*** (df = 3; 2497)	330.586*** (df = 4; 2496)	58.550*** (df = 24; 2445)	30.672*** (df = 51; 2418)

Additional controls: sex, high school type, high school final mark, standardized credits, dummies for year of enrolment.

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

The interaction term in column 4 show that treated students earn on average half additional point and the estimate is statistically significant at 10 percent level. In this specification, we include students who perhaps attended the classes, participated and benefited from online discussions but simply did not fill in the survey. Moreover, we also include students who maybe did not attend classes or use the Facebook page, because

they did not know about it or are enrolled in previous academic years. If this is the case, this set of regressions is under-estimating the true effect of online discussions and the coefficients should therefore be considered as lower-bounds.

Table 1.8: Difference-in-difference for effect of activity on math grade using administrative data also for T_{12} , so including also students who might have not attended.

	<i>Dependent variable: math grade</i>			
	(1)	(2)	(3)	(4)
January	1.525*** (0.191)	0.477*** (0.158)	0.345** (0.162)	0.267* (0.161)
T	0.886*** (0.267)	0.272 (0.219)	-0.049 (0.265)	-2.272** (0.935)
January*T	-0.322 (0.335)	0.187 (0.274)	0.553* (0.314)	0.513* (0.312)
GPA		1.098*** (0.030)	1.105*** (0.037)	1.137*** (0.037)
Constant	21.098*** (0.167)	-4.344*** (0.707)	-1.688 (3.036)	-2.595 (2.992)
Additional controls	NO	NO	YES	YES
Course FEs	NO	NO	NO	YES
Time FEs	NO	NO	NO	YES
Observations	2,704	2,704	2,647	2,647
R ²	0.032	0.354	0.367	0.391
Adjusted R ²	0.031	0.353	0.362	0.382
F Statistic	29.763*** (df = 3; 2700)	369.544*** (df = 4; 2699)	66.231*** (df = 23; 2623)	44.043*** (df = 38; 2608)

Additional controls: sex, high school type, high school final mark, standardized credits, dummies for year of enrolment.

Note:

*p<0.1; **p<0.05; ***p<0.01

A possible threat to the validity of the causal interpretation of the difference-in-difference estimates is the following. In the design of the quasi-natural experiment we use the assumption that, if online discussions had an effect on performance, it is only for

students who passed the exam within the January session, and use the following sessions as a post-experimental period. This basically means assuming a *short-run* effect. It could instead be that online discussions also have a *long-run* effect and also students who took longer to pass the exam still benefited from the Facebook page. We believe this is a minor issue and, given the difference-in-difference strategy, if this is the case we would be under-estimating the true effect of online discussions.

A second possible danger to the validity of our estimates is that we cannot exclude that also students in the control group visualized and potentially benefited from the Facebook page, perhaps hearing about it from their friends and colleagues in the treated group. Anyone with internet access could visualize the page with no need to be a Facebook user or to register on the page. However, if this actually occurred, and the effect of Facebook is positive, also in this case we are under-estimating the true effect and our estimates should be considered as a lower bound for the effect of Facebook activity on performance.

1.7 Concluding remarks

In this paper we assessed an experiment that consisted in providing online support to live lectures of a traditional mathematics course for undergraduate students. We exploit a newly constructed dataset obtained merging information collected with a survey and university records. The data provides us two cohorts of comparable students, but only one group has been exposed to the online activity on Facebook. This design allows us to use a quasi-experimental approach, overcoming some of the almost unavoidable concerns related to selection bias. As a result we are able to assess the effect of Facebook use on students' math grades by using a difference-in-difference approach. To the best of our knowledge, this is the first paper that uses a quasi-experimental design to assess the effect of Facebook use on academic performance.

The students greatly appreciated the opportunity to have such support but it is more interesting to appraise the effects on grades, which are important (albeit non-unique) indicators of students' performance. Regression models on the treated group only show that, after controlling for skill as measured by the GPA and other variables, a few measures of activity (but not all) are positively correlated with the grade. In particular, students declaring that the online support was very useful (10/10 on a Likert scale) indeed get 1 additional point on average and the frequency of visualization monotonically increases the probability to have a passing grade.

Taking advantage of the fact that parallel and formally identical courses were offered in another town at the same university, we estimated a difference-in-difference model and

compare the treated group with similar students with no availability of the online support page. This procedure has the capability to net out some possible biases in the samples and reduce self-selection effects. The main results is that the members of the treated group gain on average about one extra point (out of 30) in the math exam, when compared to other students of the same university who appear to be similar in all dimensions but took the math exam without the Facebook page.

Our results also suggest a possible differential effect due to students' individual attitude. Active students who read and oftentimes wrote posts performed better than the students who mainly read the content, but the difference is small and not statistically significant. However, the fact that the group to whom they belong as a whole has better performances suggests that the page might be more effective for under-performing students, that can benefit from interaction and information posted and shared by more active and skilled colleagues.

Other studies have shown that several alternative teaching approaches using, say, videos or blended lectures or other technical and pedagogical devices, do not change much the performance of the students as far as grades are concerned. We feel that our work is in line with such previous findings: even if we do find a significant effect, from the practical point of view, its magnitude is limited and close to $1/30$. This is less than the gap needed to move from, say, B to B+, using another familiar grading scale. In retrospect we believe this is a reasonable amount for the net improvement that can possibly be gained from activity online, as the bulk of competencies and comprehension still is (and always was) due to hard, individual and mainly *offline*, work to understand, elaborate and digest reading material and assigned problems or tasks.

We are well aware that these results should be taken as suggestive rather than conclusive and this paper could be improved in several directions. Although we are confident that we got a representative sample of the overall population of students, more research on larger representative samples would allow us to increase the external validity of results and dig into heterogeneous effects. Moreover, these results ultimately concern undergraduate economics students taking a mathematics course, and caution is needed to generalize immediately to different types of students and studied subjects.

Appendix

1.A Descriptive statistics

Figure 1.A.1: Comparison between observable characteristics in the treated and the control group used in the difference-in-difference.

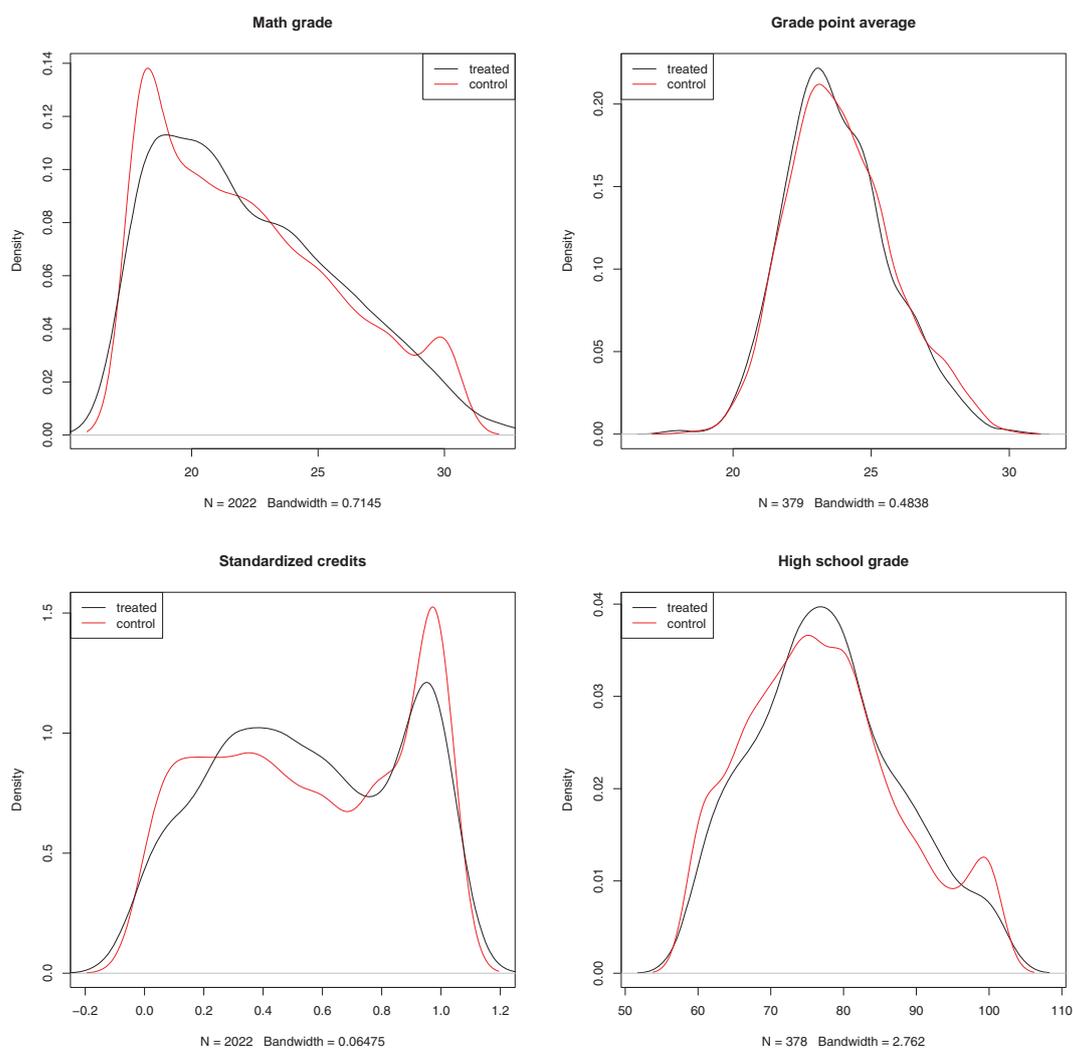


Table 1.A.1: Comparison between math grade and GPA in the treated and the control subgroups.

	T/C, January	Median	Mean
	1, 1	24.00	23.82
math	1, 0	21.00	22.02
grade	0,1	22.00	22.62
	0,0	20.00	21.08
GPA	1, 1	23.94	24.13
	1, 0	23.43	23.77
	0,1	23.95	24.14
	0,0	23.09	23.19

Table 1.A.2: Comparison between sex and high school final mark in the T and C groups.

	\bar{x}_T	\bar{x}_C	p-value
professionale	0.013	0.009	0.442
magistrale	0.020	0.014	0.362
tecnico	0.150	0.108	0.013
liceo	0.419	0.554	0.000
straniero	0.013	0.014	0.807
altro	0.384	0.300	0.000
sexM	0.36	0.501	0.000
sexF	0.64	0.499	0.000

Table 1.A.3: Comparison between teaching evaluations.

	math		all	
	T	C	T	C
2011	3.57 (4/30)	3.31	3.15	3.18
2012	3.67 (2/38)	3.33	3.22	3.25

Table 1.A.4: Comparison between the 2011 and 2012 sample, administrative data. Column 3 is the p-value of a *t-test*, where the null hypothesis is $\bar{x}_{2011} = \bar{x}_{2012}$.

	\bar{x}_{2011}	\bar{x}_{2012}	p-value
math grade	23.49	24.04	0.36
GPA	23.98	24.15	0.55
credits std	0.53	0.54	0.77
high school	80.27	79.15	0.47

Table 1.A.5: Comparison between the 2011 and the 2012 survey, 217+138=236 students. Column 3 is the p-value of a *t-test*, where the null hypothesis is $\bar{x}_{2011} = \bar{x}_{2012}$.

variable	\bar{x}_{2011}	\bar{x}_{2012}	p-value
sexF	0.66	0.65	0.90
umodatt	0.30	0.32	0.77
unmatt	0.66	0.64	0.79
uleasta	0.02	0.03	0.74
umosta	0.01	0.01	0.54
useful1	0.00	0.01	0.32
useful6	0.01	0.01	0.96
useful7	0.05	0.07	0.32
useful8	0.18	0.25	0.17
useful9	0.23	0.24	0.77
useful10	0.53	0.42	0.04
vexam	0.03	0.04	0.44
vfewtweek	0.36	0.22	0.01
vonced	0.45	0.39	0.26
vseved	0.16	0.33	0.00
vnotif	0.00	0.01	0.16

1.B Additional material

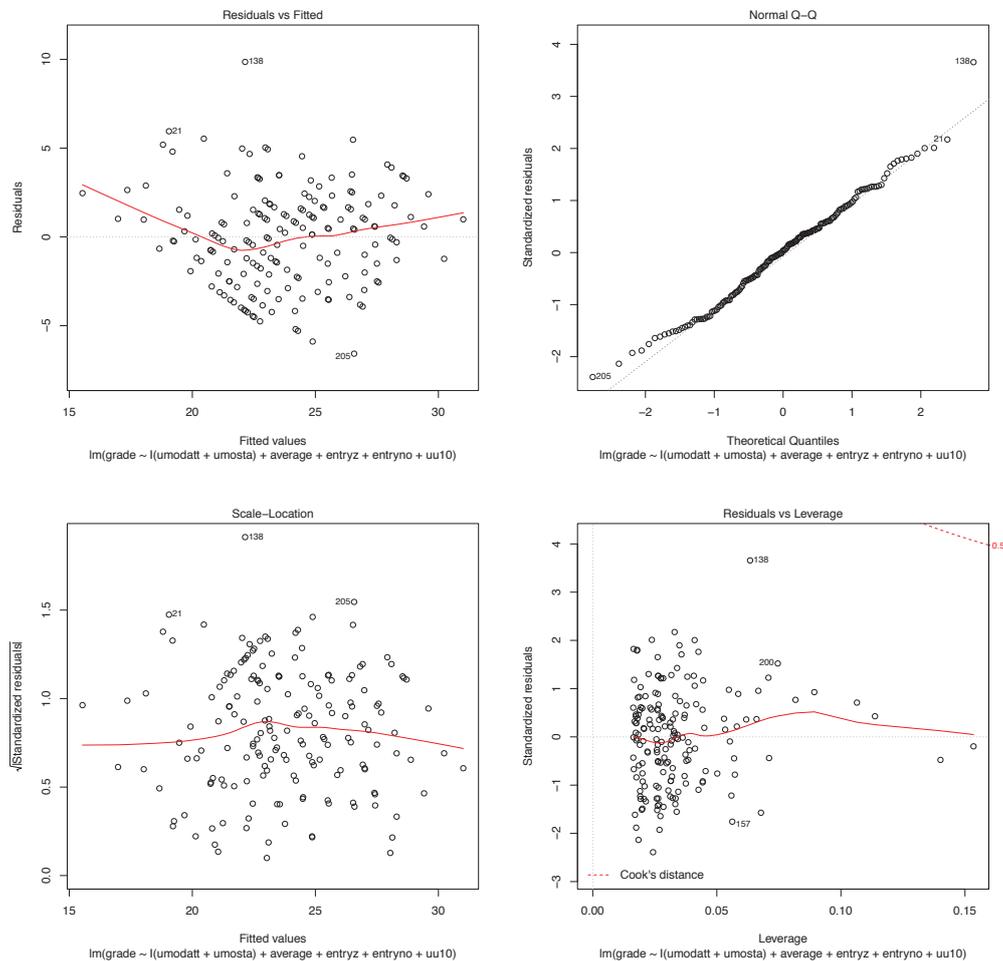


Figure 1.B.1: Residual analysis of the OLS model.

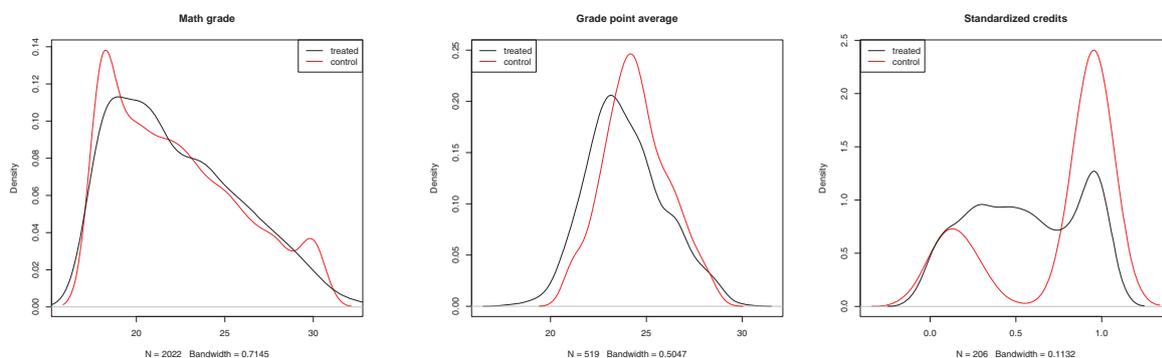


Figure 1.B.2: Comparison between treated and control group observable characteristics.

Stai visualizzando le scuole di **Indirizzo classico** in un raggio di **30Km da Venezia**

Denominazione	Ente	Prov.	Comune	Posizione	Indice FGA (1)	Media dei Voti (1)	Crediti Ottenuti (1)
MAJORANA - CORNER	●	VE	MIRANO	1	84.67	28.71	80.04
ANTONIO CANOVA	●	TV	TREVISO	2	81.67	28.41	76.56
GIUSEPPE VERONESE	●	VE	CHIOGGIA	3	79.49	27.11	83.02
BRUNO - FRANCHETTI	●	VE	VENEZIA	4	79.39	28.11	74.47
EUGENIO MONTALE	●	VE	SAN DONA' DI PIAVE	5	78.66	27.49	78.23
MARCO FOSCARINI	●	VE	VENEZIA	6	76.84	27.66	73.15
MARCO POLO	●	VE	VENEZIA	7	76.44	27.86	70.66
GIOVANNI PAOLO I	○	VE	VENEZIA	8	59.67	26.4	49.33

Stai visualizzando le scuole di **Indirizzo scientifico** in un raggio di **30Km da Venezia**

Denominazione	Ente	Prov.	Comune	Posizione	Indice FGA (1)	Media dei Voti (1)	Crediti Ottenuti (1)
GALILEO GALILEI	●	VE	SAN DONA' DI PIAVE	1	89.85	28.67	90.72
LEONARDO DA VINCI	●	TV	TREVISO	2	89.16	28.96	86.92
GALILEO GALILEI	●	VE	DOLO	3	84.82	27.92	86.92
MAJORANA - CORNER	●	VE	MIRANO	4	84.76	27.95	86.57
UGO MORIN	●	VE	VENEZIA	5	84.58	28.12	84.79
ALBERT EINSTEIN	●	PD	PIOVE DI SACCO	6	82.62	27.66	84.68
GIUSEPPE BERTO	●	TV	MOGLIANO VENETO	7	80.56	27.74	79.94
BRUNO - FRANCHETTI	●	VE	VENEZIA	8	79.26	27.61	78.39
DUCA DEGLI ABRUZZI	●	TV	TREVISO	9	78.62	26.64	85.18
BENEDETTI - TOMMASEO	●	VE	VENEZIA	10	75.65	27.26	74.06

Stai visualizzando le scuole di **Ind. Scienze Umane (ex Socio-Psico-Ped.)** in un raggio di **30Km da Venezia**

Denominazione	Ente	Prov.	Comune	Posizione	Indice FGA (1)	Media dei Voti (1)	Crediti Ottenuti (1)
GALILEO GALILEI	●	VE	DOLO	1	75.08	26.97	75.38
ALBERT EINSTEIN	●	PD	PIOVE DI SACCO	2	69.04	26.06	70.87
GIUSEPPE VERONESE	●	VE	CHIOGGIA	3	67.75	25.42	73.62
GIUSEPPE BERTO	●	TV	MOGLIANO VENETO	4	67.07	26.91	59.83
DUCA DEGLI ABRUZZI	●	TV	TREVISO	5	63.75	25.55	64.57
LUIGI STEFANINI	●	VE	VENEZIA	6	61.63	25.13	63.79
BENEDETTI - TOMMASEO	●	VE	VENEZIA	7	39.97	24.53	25.49

Tesi di dottorato "Essays on Labour Economics and Education"
di FANO SHIRA

discussa presso Università Commerciale Luigi Bocconi-Milano nell'anno 2016

La tesi è tutelata dalla normativa sul diritto d'autore (Legge 22 aprile 1941, n.633 e successive integrazioni e modifiche).

Sono comunque fatti salvi i diritti dell'università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte.

Stai visualizzando le scuole di **Indirizzo linguistico** in un raggio di **30Km** da **Venezia**

Denominazione	Ente	Prov.	Comune	Posizione	Indice FGA (i)	Media dei Voti (i)	Crediti Ottenuti (i)
ANTONIO CANOVA	●	TV	TREVISO	1	85.53	28.41	84.31
MAJORANA - CORNER	●	VE	MIRANO	2	80.51	27.28	83.66
ALBERT EINSTEIN	●	PD	PIOVE DI SACCO	3	74	26.95	73.38
LUIGI STEFANINI	●	VE	VENEZIA	4	72.45	26.56	73.51
DUCA DEGLI ABRUZZI	●	TV	TREVISO	5	71.5	26.37	73.21
BENEDETTI - TOMMASEO	●	VE	VENEZIA	6	63.81	25.73	63.16
SANTA CATERINA DA SIENA	○	VE	VENEZIA	7	63.7	26.5	56.52
SAN LUIGI	○	VE	ERACLEA	8	43.87	23.48	42.03
GALILEO GALILEI	○	TV	TREVISO	9	40.23	23.82	31.89

Stai visualizzando le scuole di **Ind. tecnico - sett. economico** in un raggio di **30Km** da **Venezia**

Denominazione	Ente	Prov.	Comune	Posizione	Indice FGA (i)	Media dei Voti (i)	Crediti Ottenuti (i)
RICCATI - LUZZATI	●	TV	TREVISO	1	69.3	25.82	73.37
GIUSEPPE MAZZOTTI	●	TV	TREVISO	2	67.59	26.07	67.88
8 MARZO - LORENZ	●	VE	MIRANO	3	66.98	25.62	70.45
LEON BATTISTA ALBERTI	●	VE	SAN DONA' DI PIAVE	4	65.21	25.46	68.19
MARIA LAZZARI	●	VE	DOLO	5	62.6	25.5	62.66
ANDREA GRITTI	●	VE	VENEZIA	6	61.78	25.37	62.07
LUZZATTI - GRAMSCI	●	VE	VENEZIA	7	60.87	25.84	56.34
ENRICO DE NICOLA	●	PD	PIOVE DI SACCO	8	60.06	25.17	60.34
FRANCESCO ALGAROTTI	●	VE	VENEZIA	9	57	25.48	51.66
CESTARI - RIGHI	●	VE	CHIOGGIA	10	48.33	24.34	43.81

Stai visualizzando le scuole di **Ind. tecnico - sett. tecnologico** in un raggio di **30Km** da **Venezia**

Denominazione	Ente	Prov.	Comune	Posizione	Indice FGA (i)	Media dei Voti (i)	Crediti Ottenuti (i)
ANDREA PALLADIO	●	TV	TREVISO	1	71.01	26.6	70.31
CARLO ZUCCANTE	●	VE	VENEZIA	2	69.96	26.68	67.54
ENRICO FERMI	●	TV	TREVISO	3	69.06	26.77	64.98
VITO VOLTERRA	●	VE	SAN DONA' DI PIAVE	4	68.27	26.18	68.36
8 MARZO - LORENZ	●	VE	MIRANO	5	65.32	26.01	63.89
SCARPA - MATTEI	●	VE	SAN DONA' DI PIAVE	6	63.87	26.1	60.23
PRIMO LEVI	●	VE	MIRANO	7	61.2	25.03	63.78
VENDRAMIN CORNER	●	VE	VENEZIA	8	50.86	24.48	47.69
CESTARI - RIGHI	●	VE	CHIOGGIA	9	45.92	24.35	38.88

Tesi di dottorato "Essays on Labour Economics and Education"
di FANO SHIRA

discussa presso Università Commerciale Luigi Bocconi-Milano nell'anno 2016

La tesi è tutelata dalla normativa sul diritto d'autore (Legge 22 aprile 1941, n.633 e successive integrazioni e modifiche).

Sono comunque fatti salvi i diritti dell'università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte.

Table 1.A.6: Summary statistics, N=236 observations (102+134).

Statistic	N	Mean	St. Dev.	Min	Max
math grade	176	23.784	3.931	18	32
pass	236	0.746	0.436	0	1
wave	236	0.568	0.496	0	1
sexF	236	0.674	0.470	0	1
GPA	236	24.078	2.144	19.000	29.000
credits	236	109.398	42.822	12	174
creditsstd	236	0.540	0.321	0.000	1.000
hsno	236	0.051	0.220	0	1
hsz	236	75.602	20.741	60	100
entryno	236	0.725	0.448	0	1
entryz	236	5.326	9.449	7.000	36.000
umodatt	236	0.309	0.463	0	1
unmatt	236	0.661	0.474	0	1
uleasta	236	0.017	0.129	0	1
umosta	236	0.013	0.112	0	1
useful1	236	0.004	0.065	0	1
useful6	236	0.008	0.092	0	1
useful7	236	0.059	0.237	0	1
useful8	236	0.169	0.376	0	1
useful9	236	0.250	0.434	0	1
useful10	236	0.508	0.501	0	1
vexam	236	0.030	0.170	0	1
vfewtweek	236	0.233	0.424	0	1
vonced	236	0.466	0.500	0	1
vseved	236	0.263	0.441	0	1
vnotif	236	0.008	0.092	0	1
u1p	236	0.153	0.360	0	1
u2p	236	0.339	0.474	0	1
nodiff	236	0.038	0.192	0	1
udontkn	236	0.089	0.285	0	1
umore2	236	0.381	0.487	0	1
fbtime	134	54.403	40.248	5	200
matematestime	133	16.511	15.989	1	100
examtime	134	142.463	102.859	2	480
lfbtime	134	3.723	0.773	1.609	5.298
lmatematestime	133	2.464	0.842	0.000	4.605
lexamtime	134	4.631	0.943	0.693	6.174
fra010	134	0.313	0.466	0	1
fra1020	134	0.299	0.459	0	1
fra2050	134	0.187	0.391	0	1
piu50	134	0.172	0.378	0	1
zero	134	0.030	0.171	0	1

Table 1.B.1: OLS Regression Results

	<i>Dependent variable: math grade</i>				
	(1)	(2)	(3)	(4)	(5)
activity	0.484 (0.484)	0.458 (0.475)	0.366 (0.462)	0.302 (0.460)	0.283 (0.458)
useful8	0.767 (1.132)	0.965 (1.082)	0.958 (1.081)		
useful9	-0.272 (1.071)	-0.148 (1.025)	-0.143 (1.024)		
useful10	1.082 (1.048)	1.178 (0.981)	1.186 (0.980)	0.961** (0.435)	0.985** (0.431)
GPA	1.322*** (0.135)	1.325*** (0.130)	1.292*** (0.123)	1.284*** (0.123)	1.217*** (0.102)
creditsstd	-0.736 (0.857)	-0.686 (0.819)			
entryz	0.219*** (0.061)	0.216*** (0.060)	0.214*** (0.060)	0.212*** (0.060)	0.218*** (0.060)
entryno	4.039*** (1.415)	3.926*** (1.325)	3.880*** (1.322)	3.912*** (1.325)	4.107*** (1.304)
hsz	-0.020 (0.025)	-0.021 (0.024)	-0.025 (0.023)	-0.022 (0.023)	
hsno	-1.417 (2.201)	-1.499 (2.126)	-1.737 (2.105)	-1.638 (2.108)	
vfewtweek	1.267 (1.744)				
vonced	1.033 (1.699)				
vseved	0.999 (1.741)				
Constant	-12.540*** (3.415)	-11.494*** (2.930)	-10.743*** (2.786)	-10.504*** (2.703)	-10.864*** (2.669)
year FEs	YES	YES	YES	YES	YES
Observations	176	176	176	176	176
R ²	0.527	0.525	0.523	0.515	0.513
Adjusted R ²	0.486	0.496	0.497	0.495	0.498
F Statistic	12.801*** (df = 14; 161)	18.231*** (df = 10; 165)	20.215*** (df = 9; 166)	25.515*** (df = 7; 168)	35.759*** (df = 5; 170)

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 1.B.2: Estimates of the selection equation of the Heckman model, 236 observations.

<i>Dependent variable: pass (1 if grade \geq 18)</i>				
	Estimate	Std. Error	t value	$Pr(> t)$
(Intercept)	-0.19	0.17	-1.13	0.26
creditsstd	1.77	0.31	5.64	5.06e-08 ***
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01			

Table 1.B.3: Heterogeneous effects of activity on Facebook use

	<i>Dependent variable: math grade</i>	
	(1)	(2)
activity	0.154 (0.818)	0.470 (0.600)
sex	-0.411 (0.555)	
low achiever		-0.578 (0.801)
GPA	1.234*** (0.104)	1.087*** (0.168)
useful10	0.989** (0.432)	0.959** (0.433)
entryz	0.217*** (0.060)	0.218*** (0.061)
entryno	4.071*** (1.311)	4.244*** (1.318)
activity*sex	0.197 (0.982)	
activity*low achiever		-0.553 (0.964)
Constant	-10.960*** (2.684)	-7.526* (4.363)
Observations	176	176
R ²	0.514	0.517
Adjusted R ²	0.494	0.497
Residual Std. Error (df = 168)	2.796	2.787
F Statistic (df = 7; 168)	25.421***	25.721***
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

Table 1.B.4: Effect of visualizations on passing probability

	<i>Dependent variable:</i>			
	pass			
	<i>logistic</i>	<i>lpm</i>	<i>logistic</i>	<i>lpm</i>
	(1)	(2)	(3)	(4)
GPA	0.440*** (0.090)	0.068*** (0.012)	0.262** (0.118)	0.038** (0.016)
vfewtweek	1.487* (0.821)	0.324** (0.146)	1.278 (0.907)	0.287* (0.151)
vonced	1.673** (0.791)	0.356** (0.141)	1.494* (0.870)	0.298** (0.145)
vseved	2.084** (0.829)	0.418*** (0.145)	1.834** (0.899)	0.352** (0.149)
Constant	-10.962*** (2.239)	-1.245*** (0.317)	-10.569*** (2.955)	-0.998*** (0.377)
Additional controls	NO	NO	YES	YES
time FEs	NO	NO	YES	YES
Observations	236	236	236	236
R ²		0.151		0.272
Adjusted R ²		0.136		0.208
Log Likelihood	-114.890		-97.286	
Akaike Inf. Crit.	239.780		234.572	
Residual Std. Error		0.406 (df = 231)		0.388 (df = 216)
F Statistic		10.265*** (df = 4; 231)		4.255*** (df = 19; 216)

Additional controls include: sex, uu8, uu9, uu10, activity, wave, creditsstd,entryz, entryno, hsz, hsno, additional2p, additionalnodiff, additionaldontkn, additionalmore2

Note:

*p<0.1; **p<0.05; ***p<0.01

Tesi di dottorato "Essays on Labour Economics and Education"
di FANO SHIRA

discussa presso Università Commerciale Luigi Bocconi-Milano nell'anno 2016

La tesi è tutelata dalla normativa sul diritto d'autore(Legge 22 aprile 1941, n.633 e successive integrazioni e modifiche).

Sono comunque fatti salvi i diritti dell'università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte.

Table 1.B.5: Activity and math grade, alternative activity measures.

	<i>Dependent variable: math grade</i>			
	(1)	(2)	(3)	(4)
average	1.257*** (0.104)	1.256*** (0.105)	1.128*** (0.124)	1.282*** (0.104)
I(umodatt + umosta)	0.521 (0.473)			
vfewtweek		1.285 (1.755)		
vonced		1.224 (1.723)		
vseved		1.405 (1.741)		
lexamtime			-0.398 (0.267)	
lmatematestime			-0.606 (0.407)	
lfbtime			-0.136 (0.437)	
uu10				0.935** (0.438)
Constant	-7.191*** (2.565)	-8.276*** (3.024)	-0.107 (3.700)	-8.140*** (2.589)
Observations	176	176	93	176
R ²	0.459	0.458	0.513	0.470
Adjusted R ²	0.453	0.445	0.491	0.464
F Statistic	73.532*** (df = 2; 173)	36.108*** (df = 4; 171)	23.193*** (df = 4; 88)	76.598*** (df = 2; 173)

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 1.B.6: Activity and passing probability

	<i>Dependent variable: pass (1 if math grade \geq 18)</i>			
	(1)	(2)	(3)	(4)
average	0.251*** (0.049)	0.253*** (0.050)	0.344*** (0.069)	0.252*** (0.049)
I(umodatt + umosta)	-0.091 (0.197)			
vfewtweek		0.919* (0.489)		
vonced		1.041** (0.471)		
vseved		1.265*** (0.490)		
lexamtime			-0.078 (0.148)	
lmatematestime			0.210 (0.189)	
lfbtime			-0.003 (0.196)	
uu10				0.319* (0.188)
Constant	-5.262*** (1.160)	-6.357*** (1.252)	-7.751*** (1.889)	-5.473*** (1.161)
Observations	236	236	133	236
Log Likelihood	-118.401	-114.775	-64.328	-117.062
Akaike Inf. Crit.	242.802	239.551	138.656	240.124

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 1.B.7: Comparison between the treated group T and the control group C .

	\bar{x}_T	\bar{x}_C	p-value
math grade	22.68	21.84	0.06
GPA	23.93	24.46	0.00
credits std	0.56	0.73	0.00
high school	77.847	77.269	0.306
professionale	0.013	0.013	0.977
magistrale	0.021	0.007	0.064
tecnico	0.146	0.232	0.003
liceo	0.420	0.536	0.001
straniero	0.012	0.010	0.826
altro	0.387	0.199	0.000
sexM	0.355	0.248	0.001
sexF	0.645	0.748	0.002

Table 1.B.8: Comparison between the treated group T and the control group C .

	T/C, january	Median	Mean
	1, 1	24.00	23.82
math	1, 0	21.00	22.02
grade	0, 1	22.00	21.99
	0, 0	21.00	21.69
	1, 1	23.94	24.13
GPA	1, 0	23.43	23.77
	0,1	24.77	24.91
	0,0	24.05	24.00

Table 1.B.9: Difference-in-difference estimates, students of previous years as control group.

<i>Dependent variable: math grade</i>			
	(1)	(2)	(3)
January	0.015 (0.399)	-0.773** (0.331)	-0.810** (0.334)
T	0.151 (0.351)	0.639** (0.290)	0.569* (0.302)
January*T	1.781*** (0.521)	1.733*** (0.428)	1.779*** (0.433)
GPA		1.071*** (0.056)	1.126*** (0.066)
Constant	21.870*** (0.287)	-4.085*** (1.380)	-4.165*** (1.461)
Additional controls	NO	NO	YES
Observations	764	764	761
R ²	0.048	0.357	0.363
Adjusted R ²	0.045	0.354	0.353
F Statistic	12.858*** (df = 3; 760)	105.335*** (df = 4; 759)	35.499*** (df = 12; 748)

Note: *p<0.1; **p<0.05; ***p<0.01

Additional controls: high school mark, type of high school, credits std

Chapter 2

Dual labor markets, workers' effort and labor productivity

Abstract

The simultaneous increase in the use of temporary contracts and the productivity slow-down recently experienced in some OECD countries, fostered a growing interest in analysing the link between these phenomenon.

In this paper we study the effect of the use of temporary contracts on workers' incentives and, as a consequence, on labor productivity. We implement an agent-based model where workers interact in the labor market and compete for permanent contracts. Workers face the following trade-off: exerting high effort is a costly investment but it increases the probability that their contract is converted into permanent. Workers choose how much effort to exert in production and, using reinforcement learning, they update their strategies on the basis of past experience.

The main result is that optimal effort strategies depend on the share of available permanent contracts. When the share is low, workers do not bet on their conversion and supply low effort. As the share increases workers exert higher effort but, when it is too high, they have the incentive to shirk since they are too confident of being confirmed. As a consequence, the relationship between the share of permanent contracts and labor productivity has an inverted-U-shape.

Keywords: temporary employment, workers' effort, reinforcement learning.

2.1 Introduction

In the recent years a large number of reforms has been implemented in different European labor market institutions. As shown in Table 2.1, most of them are *two-tier reforms*, meaning they have an effect only on part of the population, also at the steady state. These reforms contribute to create what in the literature is called a *dual labor market*,

Reform area	Two-tier	Complete	Total per row	Of which two-tier
EPL	103	96	199	52%
UB	116	137	253	46%
AP	155	87	242	64%
ECI	74	50	124	60%
ER	49	16	65	75%

Table 2.1: Summary of reforms implemented between 1980 to 2007 in employment protection legislation (EPL), unemployment benefits (UB), active labor market programs (AP), employment conditional incentives (ECI) and early retirement plans (ER); source Boeri (2011).

featuring the coexistence of two main types of contracts: permanent contracts, open-ended contracts protected with high firing costs, and temporary contracts, with limited duration and low or no firing costs.

The motivation of this paper builds on the observation of two empirical facts: on one hand, an increase in the use of temporary contracts and, on the other, a slowdown in labor productivity in some OECD countries. Figure 2.1 depicts the increasing trend in the use of temporary contracts in the period 1980–2015 (*left*) and the slowdown in labor productivity in Italy compared to other OECD countries (*right*).

In Italy, for example, the sharp increase in the share of temporary contracts is due to a number of reforms that increased the possibility of using this type of agreements; see Cappellari et al. (2012).¹ Since then, temporary contracts are typically used for many different reasons: screening purposes, temporarily fill-in for staff who are absent or on leave, or to accommodate fluctuations in demand; in many cases employers also save in labor costs and social security benefits.

¹Before 2001 the law regulating fixed-term contracts had a very specific list of circumstances under which a firm could use a fixed term contract but the law 368/2001 eased restrictions on fixed term contracts, introducing a single general requirement. They could be implemented “...for reasons of a technical, organizational, production or replacement nature”.

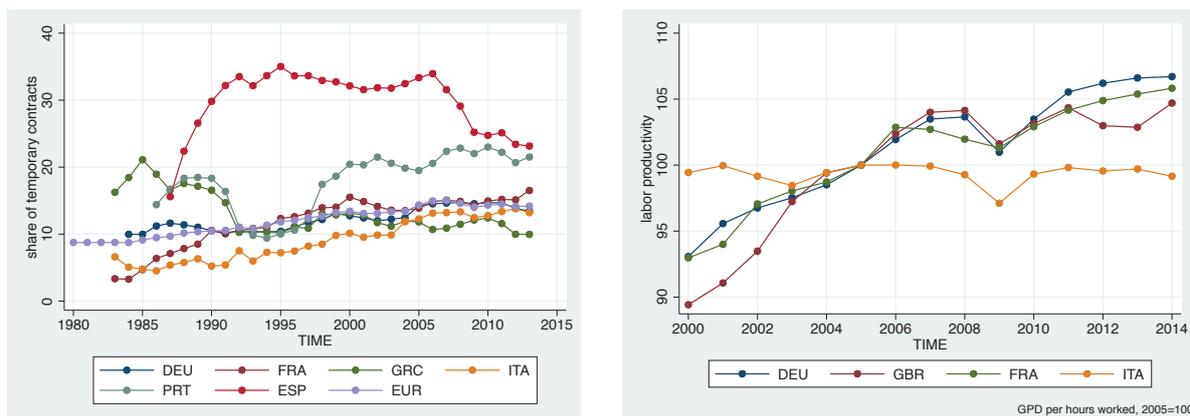


Figure 2.1: Share of temporary contract for selected OECD countries in the period 1980–2015 (*left*) and time series of labor productivity in the years 2000–2014 (*right*). Data source is OECD database STAN: <http://www.oecd.org/sti/ind/stanstructuralanalysisdatabase.htm>

The aim of this paper is to study the link between the use of temporary employment and labor productivity. The channel we have in mind is that temporary contracts have an effect on workers' incentives, and in particular on their willingness to exert high effort and, consequently, this affects firms' labor productivity. We implement an agent-based model where workers and firms interact in a dual labor market.² In the model, temporary workers compete for a limited number of permanent contracts and they face the following *trade-off*: exerting high effort is costly but it increases the chances of obtaining a permanent contract. In this environment, workers choose how much effort to exert in the production process and update their strategies on the basis of past experience; agents use reinforcement learning as a learning algorithm, see Roth and Erev (1995).

The use of an agent-based model has the advantage that it explicitly and easily allows us to take into account the interaction among workers competing for permanent contracts, where the probability for a worker of obtaining a permanent contract depends not only on his effort, but also on effort exerted by the other temporary workers in the firm. To the best of our knowledge, this channel has previously been neglected in the literature.

This paper is related to different strands of literature: empirical papers investigating the effect of temporary employment on labor market outcomes, and in particular on productivity; studies of workers' behaviour and incentives under different contractual schemes and agent-based models of the labor market.

Given the increase of temporary employment several empirical studies have addressed

²For an introduction to complexity and agent-based models see Miller et al. (2008) and Batten (2000) <http://www2.econ.iastate.edu/classes/econ308/tesfatsion/Batten.EntireBook.pdf>.

this topic, looking for example at the effects on unemployment, turnover, workers' careers (Gagliarducci (2005) and Güell and Petrongolo (2007)) and on labor market performance (Blanchard and Landier (2002) and Cahuc and Postel-Vinay (2002)). An increasing number of contributions focuses on the effect of temporary employment on productivity; different channels have been proposed and most of the papers find that temporary contracts have a negative effect on firm productivity.

A first channel is given by a reduction in the firm's propensity to innovate: firms might instead choose to maintain labor intensive production and aim at cutting cost, see for example Michie and Sheehan (1999). A second channel through which temporary employment can affect productivity is training and human capital accumulation; employers may be reluctant to invest in human capital of fixed term workers, because the period in which the firm would benefit from the investment is too short. Moreover, temporary workers themselves may not be willing to acquire firm specific competences, if they don't feel committed to a firm, see Albert et al. (2005) and Kleinknecht et al. (2006). A third channel indicates that temporary employment may have a negative effect on productivity because this variable is positively correlated with experience and permanent employees may have more experience than temporary ones; see Daveri and Parisi (2010).

Focusing on Italy, the boom of atypical contracts and the productivity slowdown that simultaneously occurred stimulated empirical analysis of the relationship between the two events at the firm level. Boeri and Garibaldi (2007) found a negative effect of the share of fixed-term contracts on labour productivity growth in a sample of manufacturing firms already during the period 1995–2000. More recently, Lotti and Viviano (2012) provided further evidence for the existence of a negative relationship between the use of fixed-term contracts and the lower labour productivity of manufacturing firms. Cappellari et al. (2012) investigates the effects of two reforms of temporary employment using panel data on Italian firms. Their results show that the reforms of fixed-term contracts induced a substitution of temporary employees in favour of external staff and reduced capital intensity, generating productivity losses.

Nevertheless, fixed-term contracts may also have some benefits on firms' outcomes. For example, flexible contracts allow firms to adopt the employment level to fluctuations in demand, increasing efficiency at the margin. Moreover, more flexible labor and consequent higher turnover, might favour firms' innovation and the entrance of new ideas in the firms.

While several papers look at the effect of temporary employment on productivity, less is known about the impact on workers' behaviour; in this paper we provide evidence for the latter and in particular we assess whether and under which conditions temporary

employment induces workers to exert high effort. In this respect, our results are complementary to the literature focusing on the effect of employment protection on workers' effort and the efficiency wage theory proposed by Shapiro and Stiglitz (1984) and tested by Cappelli and Chauvin (1991).

Guadalupe (2003) is the first paper that looks at behavioural responses of temporary contracts showing that they cause significantly higher accident rates; our focus is instead on workers' effort. Few empirical papers have documented the effect of temporary employment and labor market regulations on workers' effort decisions; this is mainly related to the difficulty of finding good proxies of effort, that is an unobservable variable. Three examples using respectively Swiss, Italian and Spanish data are respectively Engellandt and Riphahn (2005), Ichino and Riphahn (2005) and Dolado et al. (2012).

In the first paper, effort is proxied by unpaid overtime work and absences; the authors provide evidence that temporary workers provide higher effort than permanent employees. Their probability of working overtime exceeds that of permanently employed workers by 60%. The second paper looks at the effect of a change in the employment protection legislation regime on absenteeism, used to proxy effort. The main finding is that the number of days of absence per week increases significantly once employment protection is granted at the end of a three month probation period, therefore, highly protected contracts may induce lower effort. Finally, Dolado et al. (2012) analyses the effect of having a large gap in firing costs between permanent and temporary workers on total factor productivity at the firm level. The authors show that firms' temporary-to-permanent conversion rates and consequently temporary workers' effort decrease when the gap increases.

Differently from the previous contributions, using an agent-based model we take into account the additional competition channel faced by temporary workers competing for permanent positions in the same firms. Leading labor economists such as Freeman (1998) suggest the use of these techniques to study the interaction between workers and firms in the labor market, with the aim of replicating stylized facts and analysing the effects of specific policies (e.g. training policies, unemployment benefits etc...). According to Freeman (1998), "*...the set of new tools such as simulations, complexity analysis, neural networks and artificial agent simulated societies can be used to study labor market systems, offering a way of modelling or exploring alternative economic models through simulations that were beyond our capabilities just a few years ago*". For a recent review on agent-based models applied to labor markets see Neugart et al. (2012a). Early examples are Bergmann (1973) and Eliasson (1977), that model respectively the US and the Swedish economy, and Bergmann (1990), a simple model³ in which the author studies the effect of

³Workers are homogeneous, labor demand is exogenous, matching is random, unemployed workers

the introduction of unemployment insurance. Neugart (2008) is a recent example where the author uses an agent-based model to study the effect of introducing a training policy, that subsidizes workers' acquisition of skills, on unemployment. The paper shows that on aggregate the policy reduces unemployment, but it also has distributional consequences, as workers that do not receive the subsidy are harmed as they face higher competition. As in our model, workers use reinforcement learning to update decisions on how much to invest in acquiring additional skills. Other examples are Tassier and Menczer (2008) and Tesfatsion (2001b).

The rest of the paper is organized as follows: section 2.2 describes the characteristics of the model, section 2.3 presents the computational results, section 2.4 are robustness checks and section 2.5 summarized and concludes the paper.

2.2 The model

In the labor market there are N_W workers and N_F firms with $N_W \gg N_F$. Each worker is endowed with one unit of labor which is the only factor of production in the economy; firms supply an homogeneous good. Initially, workers are randomly assigned to firms, all firms employ the same number of workers and all vacancies are filled, therefore, labor force participation is constant. Initially, u workers are not allocated to any firm and start the period as unemployed.

We simply assume that the production function of firm j is defined as the sum of effort provided by the firms' employees:

$$Y_j = \sum_{i=1}^{q_j} e_{ij} \quad (2.1)$$

where q_j is the number of workers employed in firm j and e_{ij} is effort exerted by worker i when matched with firm j .

Two types of contracts characterize the labor market: temporary and permanent contracts; what makes the two contracts different is their duration. Workers with permanent contracts remain matched with the same firm, unless the firm is hit by an exogenous shock that destroys all permanent contracts in the firm. Workers with temporary contracts are employed for a maximum amount of time d ; during this period they remain temporary or can become permanent. After d rounds their contract ends and, if their contract has not been converted into permanent, they become unemployed.

The problem faced by workers is deciding how much effort they should exert in the production process. The strategy S_i of a worker is defined by a probability distribution

always accept an offer and wages are not modelled.

over the discretized set of effort choices $e_k \in [1, 2, \dots, 10]$ such that $\sum_{k=1}^{10} p_k = 1$ and $p_k \geq 0$ for all workers. Initially, workers do not know how much effort they should exert in the production process, therefore, all strategies are chosen with equal probability. We simulate the model under two different scenarios. In the first, we assume that workers stick to the strategy S_i also when their temporary contract is upgraded to permanent. In the second, instead, we assume that when workers are converted into permanent they do not sample an effort value from their distribution and, instead, exert a fixed level of effort \bar{e} , exogenously determined. The intuition is that, for example, when converted into permanent workers might decide to lower their effort. Unemployed workers exert 0 effort.

All workers are initially employed with temporary contracts. Once workers are allocated to firms, they sample an effort value from their distribution and production occurs.

Each firm can employ only a given fraction P of workers with permanent contracts, for example due to institutional regulations or financial constraints. If in a firm the current fraction of permanent contracts is smaller than P the conversion process takes place. We assume that firms can observe the level of effort exerted by workers with temporary contracts. Firms therefore rank temporary workers by decreasing level of effort and the top ranked temporary workers become permanent, until the share P is reached, while the others remain temporary; ties are broken randomly.

We assume the utility of a worker with a permanent contract is greater than the utility of a worker employed with a temporary contract. Therefore, workers within each firm compete for permanent contracts. Temporary workers face the following trade-off. Exerting high effort is a costly investment, but it increases the probability that their contract is converted into permanent. Workers suffer a loss of value that is strictly positive and increasing in effort, $c(e) = \alpha e^\beta$ with $\alpha > 0$ and $\beta \geq 1$. All workers simultaneously make their effort decisions, without knowledge on the level of effort exerted by the other workers in the firm.

After the conversion process takes place, workers learn their new status, permanent if their contract has been upgraded and temporary if not. The payoff of worker i is defined as:

$$\pi_i(e_i, e_{-i}) = \begin{cases} w - \alpha e_i^\beta + x_T & \text{if temporary} \\ w - \alpha e_i^\beta + x_P & \text{if permanent} \end{cases} \quad (2.2)$$

where we assume that all employed workers receive the same exogenous wage w . $X \in \{x_T, x_P\}$ is a non-monetary benefit that is different according to the type of contract and takes two values: x_P and x_T , respectively for permanent and temporary workers, with

$x_P > x_T$.⁴

In the initial stage workers choose the level of effort to exert from the set of feasible strategies with equal probability. Workers keep track of payoffs obtained with the different strategies and, as time passes, they realized that some strategies work better than others. Workers learn how much effort they should exert in the production process only when they are temporary and they are in a firm that can convert some temporary workers into permanent, to reach the share P , so after the job destruction of permanent contracts occurs. We model this as a process of individual reinforcement learning; average payoffs drive the learning process. Worker i will choose effort e_k with probability:

$$p(i, k) = \frac{\exp^{\lambda \cdot \text{avgpayoff}(i, e_k)}}{\sum_{k=1}^{10} \exp^{\lambda \cdot \text{avgpayoff}(i, e_k)}} \quad (2.3)$$

where $\lambda > 0$ determines the speed of the learning process and $\text{avgpayoff}(i, e_k)$ are average payoffs of worker i , when he played strategy e_k . In this learning process strategies that lead to relatively higher payoffs will be played with higher probability in the next rounds.⁵

Every d rounds temporary contracts end and workers separate from firms. Workers that were previously unemployed become temporary and the temporary workers that became unemployed are either randomly matched to a new firm (or the same one by chance) or remain unemployed.

Moreover, in a randomly determined order, in each round one firm is hit by a job destruction shock: permanent workers become unemployed and the vacancies are filled by previously unemployed workers, that are hired with temporary contracts. In the following round this firm will update its share of permanent contract, to reach the share P . A new round begins, with an updated allocation of workers, contracts and strategies.

In each round, the number of employed and unemployed workers remains constant, as firms fill all their vacancies, but workers reallocate across the three states (temporary, permanent and unemployed). We call period a sequence of N_F rounds, such that each firm has updated once the share of its permanent workers. After enough stages, workers learn the optimal level of effort they should exert to maximize their expected payoffs.

⁴Written in this way, the payoff is simple to interpret, but w , x_T and x_P are constant parameters therefore we could simplify the expression.

⁵Equation (2.3) is known as Gibbs-Boltzmann probability measure.

2.3 Computational results

In this section, we discuss the results of the model for a representative set of parameters, Table 2.1 summarizes our choices. When more than one possible value is given, we use boldface to denote the baseline case.

Parameter	Description	Value
N_W	number of workers	300
N_F	number of firms	30
q	number of workers per firm	10
u	number of unemployed	10
d	temporary contract maximum duration	10
γ	job destruction of permanent contracts	$1/N_F$
λ	learning parameter	{5, 10 , 20}
Δ	effort tick size $e_i \in [1, 2, \dots, 10]$	1
P	share of permanent contracts	{0.1, 0.2, ..., 1}
t	periods	500
w	wage	1
α	cost of 1 unit of effort	$U \sim [0.05, 0.15]$
β	convex cost parameter	{ 1 , 1.5, 2}
x_T	non-monetary benefit if temporary	-1
x_P	non-monetary benefit if permanent	0

Table 2.1: Description and value of the parameters used for the simulations.

We consider an economy with $N_W = 310$ workers (of which $u=10$ are unemployed) and 10 workers are allocated to each firm, hence, there are $N_F = 30$ firms in the labor market. Not all parameters are equally relevant. For example, given the setup of the model, the number of firms in the labor market is not a crucial variable in determining the models' outcomes. We assume that in a given simulation all firms can employ the same number of workers with permanent contracts. Therefore, we average across firms simply to net out sampling variation. Also the share of unemployed workers, which is very low compared to empirical evidence, is not a critical variable in the model; the batch of unemployed workers is used to replace workers who loose their job, and unemployed workers are inactive and simply wait to be hired.

We assume all workers earn an exogenous wage and normalize it to $w = 1$; the non-monetary benefit temporary workers get if they are (not) converted into permanent is set

to 0 (-1). We observe interactions among workers in the labor market for 500 periods.

In the model the workers move across different states and they can be permanent, temporary or unemployed at different times. Nevertheless, we will focus on the second group and in particular on strategies learnt by temporary workers when they are in firms that can upgrade some contracts into permanent. We start by looking at average payoffs earned by temporary workers during the reinforcement learning process and show that the algorithm converges to a steady state which is an approximation of an equilibrium. Then, we describe the strategies evolved by temporary workers at the end of the simulation.

2.3.1 Payoffs

As the learning process takes place temporary workers update their strategies and increase the probability of playing effort choices that lead to higher payoffs. We repeat a complete run of the model (500 periods) for each of the possible values of the share of permanent contracts.

Figure 2.1 depicts the time series of average payoffs earned by temporary workers, when they are in firms that can convert workers into permanent, as the learning process takes place. In each data-point all firms updated once their share of permanent contracts to reach the desired share P , therefore, each point is the average payoff of $N_W = 300$ workers.

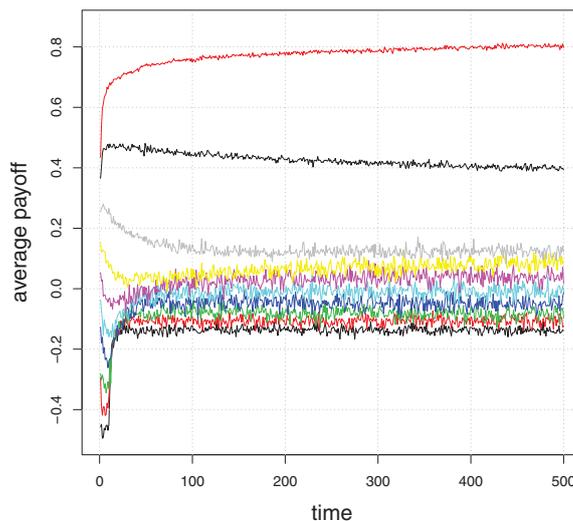


Figure 2.1: Time series of average payoffs for different values of the share of permanent contracts, from bottom to top $P = \{0.1, 0.2, \dots, 1\}$.

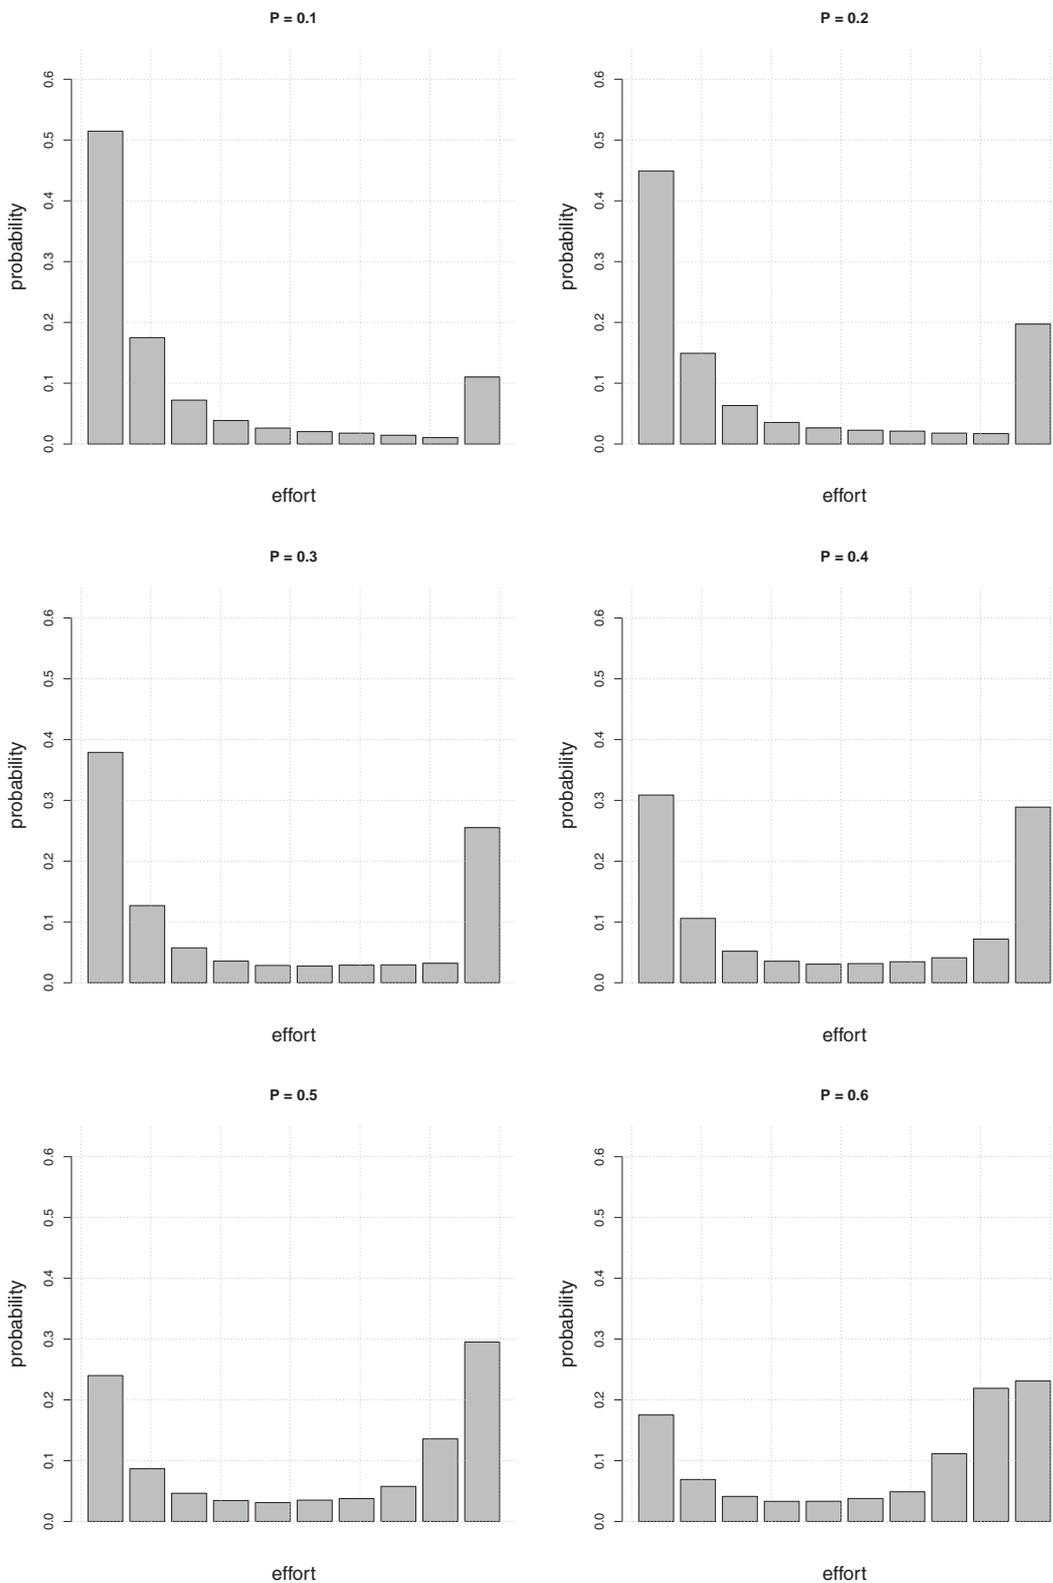
In the plot, time is on the x -axis and average payoffs of workers on the y -axis. Each time series corresponds to a simulation for a different share of permanent contracts and moving from bottom to top $P = \{0.1, 0.2, \dots, 1\}$. As expected, the higher is the share of available permanent contracts P the larger are average payoffs earned by workers. Nevertheless, in the different simulations workers have different learning patterns. For low values of P , say up to $P = 0.6$, payoffs initially decrease and then increase before converging; for higher values of P instead payoffs follow an opposite pattern, first increasing than decreasing and for $P = 1$ payoffs increase over time.

Moreover, the speed of convergence differs across simulations. For low values of the share of permanent contracts, say up to $P = 0.7$, the learning process is fast and approximately 10 to 15 updates are sufficient to reach convergence. Instead, for higher values of P the learning process is faster in the initial stages but converges slowly in a high number of iterations. While experiencing approximately 10 different temporary jobs, as in the cases $P = \{0.1, 0.2, \dots, 0.7\}$, is a high number but still of reasonable magnitude, in the cases $P = \{0.8, 0.9, 1\}$ the model does not match empirical evidence as 100 updates, or more, are needed. Recall that in this set of simulations the parameter λ determining the speed of the learning process is set to 10, whereas, using different values for the different maximization problems, might be a better choice.

2.3.2 Strategies

Figure 2.2 depicts strategies learnt by workers at the end of a simulation; each barplot corresponds to a different value of the share of permanent contracts $P = \{0.1, 0.2, \dots, 1\}$. At the end of the learning process workers evolve strategies that are very similar, but not identical both because of sampling variation and because workers have different effort costs, sampled from a uniform distribution. Therefore, we compute and plot workers' strategies averaging across all workers in the model, at the end of a simulation. The support of effort values is on the x -axis and the probability of choosing each strategy on the y -axis.

Some patterns clearly emerge. As shown by the plots, workers' optimal strategy changes as a function of the share of available permanent contracts and, in general, workers learn to play mixed strategies. For example, when $P = 0.1$ workers obtain higher payoffs exerting minimum effort and therefore play $e = 1$ with higher probability. In this case, only one worker out of ten is promoted to a permanent position, therefore, workers realize that it is not worth it to bear high effort costs.



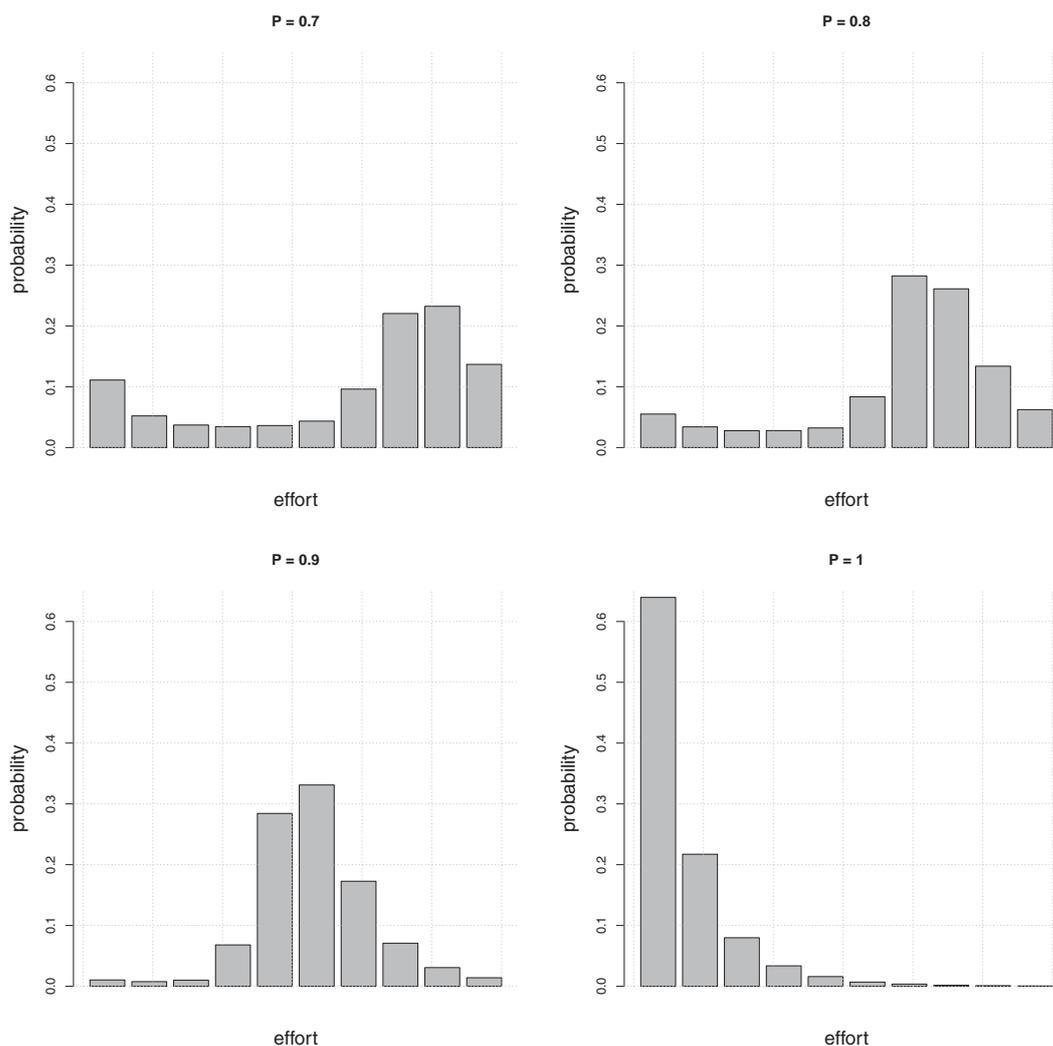


Figure 2.2: Workers' average strategy at the end of a simulation for all possible values of the share of permanent contracts in $P \in \{0.1, 0.2, \dots, 1\}$.

Nevertheless, when workers exert maximum effort chances are high that they are chosen for promotion, therefore, on average the strategy $e = 10$ is chosen with 11 percent of probability. Moving from the case $P = 0.1$ to $P = 0.2$ the probability of playing $e = 1$ decreases from 0.51 to 0.45 and the probability of playing $e = 10$ increases from 0.11 to 0.2. As the share of available permanent contracts increases workers realize that their chances of being promoted increase, therefore, they optimally decrease the probability of exerting minimum effort and increase the probability of exerting maximum effort. This is true up to the case $P = 0.5$ but, as the share of available permanent contracts further increases, something changes. Workers decrease their effort and, as P increases, the distribution shifts to the left; when $P = 1$ workers choose to exert the minimum effort level with probability 0.64.

What is going on? Why is there a tipping point after which temporary workers decrease their effort? The intuition is the following. Within each firm, temporary workers compete among each other to get permanent contracts and, when only few workers can be converted into permanent, competition induces workers to increase their effort, as the number of available promotions increases.

But, when the share of permanent contracts increases above $P = 0.5$, the pressure to compete for permanent contracts decreases as workers realize that, even exerting low effort, their contract will anyway be upgraded into permanent with high probability. In other words, temporary workers have an incentive to shirk as there is no need for them to work hard and compete for permanent contracts.

2.3.3 Effort and productivity

As seen so far, the share of available permanent contracts induces different incentives on workers' willingness to exert low/high effort and as a consequence shapes workers' strategies. We now take an aggregate approach and look at the effect of the trade-off faced by temporary workers on average effort and consequently on firm productivity.

Figure 2.3 *left* shows the expected value of effort exerted by all temporary workers in the model at the end of a simulation, it summarizes results for 10 different runs of the model, one for each of the possible values of the share of permanent contracts $P = \{0.1, 0.2, \dots, 1\}$. In the plot, the share of permanent contracts is on the x -axis and average effort on the y -axis. Each data-point is the expected value of effort from a different simulation and it is computed using the average strategies plotted in the previous figure.⁶

The main message is that the relationship between the share of permanent contracts and average effort has an inverted-U-shape. As the number of available contract upgrades increases, workers learn to exert higher effort, as this increases their chances to get a permanent contract, but, if the share of permanent contracts is too high average effort decreases. As previously discussed, temporary workers do not feel the pressure to compete for promotions as, with high probability, they will anyway get a permanent contract and therefore they decrease their effort. If temporary workers do not change their effort strategies when they become permanent, from firms' point of view it is optimal to convert 80 percent of temporary workers into permanent as at this point effort is maximized. In this case average effort is 6.87 out of 10, instead, the lowest levels of effort correspond

⁶For example in the case $P = 1$ the expected value of effort is: $E(e|P = 1) = \sum_{e=1}^{10} e \cdot p_e = [0.64 \ 0.22 \ 0.08 \ 0.03 \ 0.02 \ 0.01 \ 0 \ 0 \ 0 \ 0]' [1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7 \ 8 \ 9 \ 10] = 1.62$.

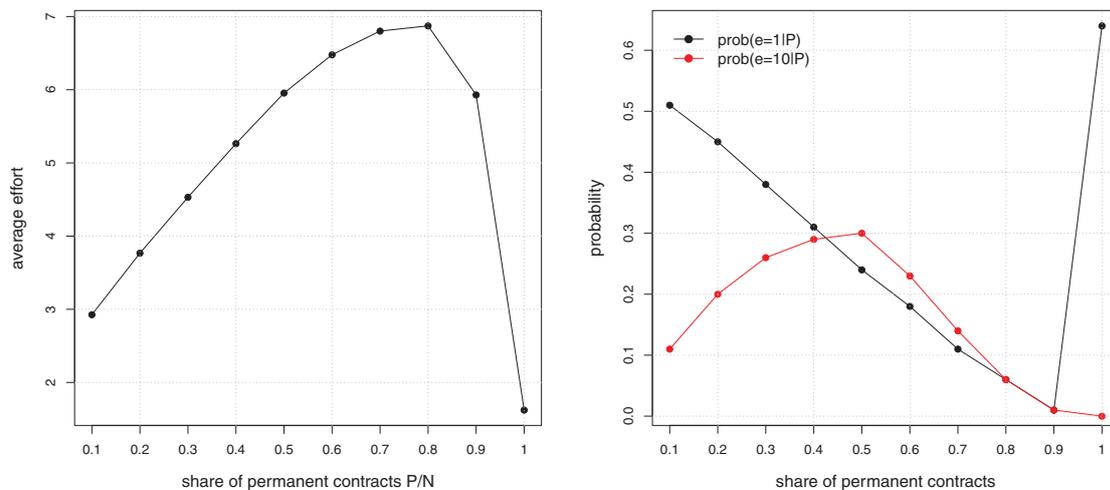


Figure 2.3: Workers' average effort (*left*) and probability of exerting minimum/maximum effort (*right*) for each value of the share of permanent contracts $P \in \{0.1, 0.2, \dots, 1\}$.

to the cases in which all workers are temporary and all workers are permanent, average effort is respectively 2.93 and 1.62. Comparing these two options the model suggests that firms would prefer to have all workers employed with temporary contracts rather than all permanent ones.

Figure 2.3 *right* conveys a similar message to the the plot on the *left*. The share of permanent contracts is on the x -axis and the probability of exerting the minimum (black line) and maximum (red line) level of effort is on the y -axis. Probabilities are averaged across all temporary workers' strategies. The probability of exerting the maximum level of effort, $e = 10$, first increases, up to the case $P = 0.5$ and then decreases. The probability is equal to 0.11, 0.3 and 0 respectively in the cases $P = \{0.1, 0.5, 1\}$. The probability of exerting the minimum level of effort, $e = 1$, follows instead an opposite patten, first decreasing then increasing.

Recall that in this labor market production of firms is simply defined as the sum of effort exerted by employed workers. Figure 2.4 *left* and *right* shows the relationship between the share of available permanent contracts and labor productivity at the firm level, simply computed as the sum of effort over the number of workers, within each firm.

The difference in the two plots is the assumption on workers' behaviour when their contract is converted into permanent. In the simulations corresponding to the *left* plot we assume permanent workers stick to the strategy they learnt when they were temporary and competing for promotion, and simply sample an effort value from their distribution

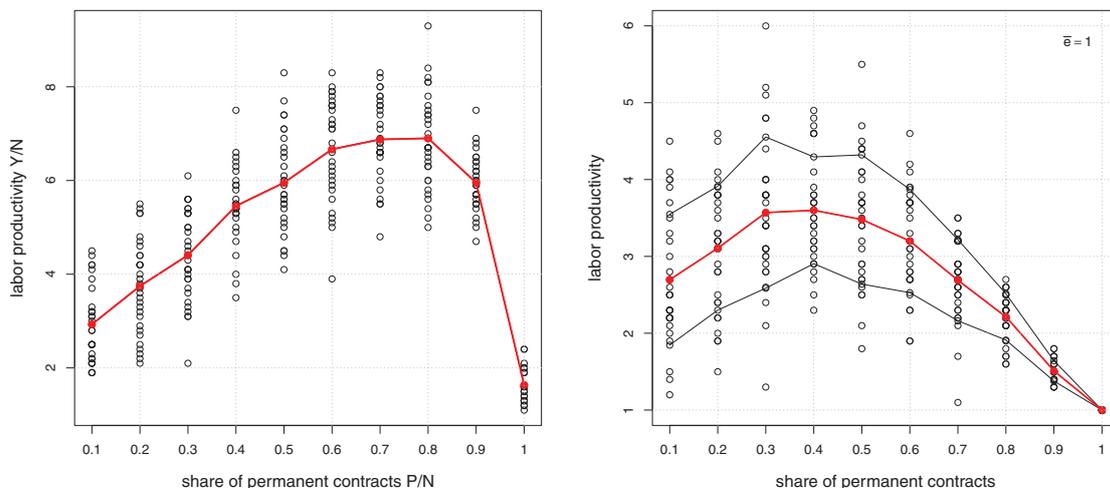


Figure 2.4: Share of permanent contracts and labor productivity under two different assumptions on permanent workers' behaviour. The red line is the mean and black is plus/minus 1 standard deviation.

also when they are permanent. On the *right* instead we assume that when workers are converted into permanent they switch to a fixed level of effort \bar{e} that is set to the minimum level $e = 1$. Each data-point represents labor productivity of one of the 30 firms in the labor market; the red line joins the average labor productivity for all values of the parameter P .

Looking first at the *left* plot, we can see there is a considerable amount of noise and different firms have different productivity values, even when they have the same share of permanent workers. Nevertheless, when considering average values the outcome is the expected one and the relationship between the share of available permanent contracts and labor productivity shows the same pattern as Figure 2.3 *left* on workers' effort, as is the model productivity is a consequence of effort decisions.⁷ Therefore if firms could observe workers' effort decisions and the conversion process was not costly, firms would optimally set the share of permanent contracts to 80 percent, inducing high effort, to maximize labor productivity.

Empirical evidence instead shows that yearly transition probabilities from fixed-term to permanent contracts are relatively small, they never exceed 50% and are as low as 12-13% in Portugal and Spain.⁸ In light of the model, this low temporary-to-permanent

⁷We obtained similar plot joining the medians.

⁸This estimate is from the European Union Survey of Income and Living Conditions (EU-SILC) in the 2004-7 period.

conversion probability may be one of the factors causing low productivity values recently observed in some OECD countries, as firms are not providing temporary workers the “right” incentives to exert high effort.

Figure 2.4 *right* summarizes the outcome of the model when we assume that workers that become permanent switch to the minimum level of effort $\bar{e} = 1$. Also in this case, the plot shows that the relationship between the share of permanent contracts and labor productivity has an inverted-U-shape, as emphasised by the red line joining average labor productivity values. Under this assumption, the model suggests that firms should convert 40 percent of temporary workers into permanent to maximize labor productivity. At this point average labor productivity is 3.6. Note that the scale on the y -axis on the two plots is different and, as expected, higher levels of labor productivity can be reached in the scenario plotted on the *left*. As the share of permanent contracts increases, more workers exert the minimum level of effort \bar{e} , decreasing labor productivity. The black lines are average labor productivity plus/minus one standard deviation. Note that, as expected, when the share of permanent contracts increases above 50 percent, the standard deviation monotonically decreases as more workers are exerting the minimum effort level.⁹

2.4 Robustness checks

In this section we test the robustness of our results with respect to several parameters used in the baseline model.

Not all parameters have the same importance in determining the models’ outcomes, and in particular the effort strategies learn by temporary workers.

We start by looking at the robustness of results with respect to the parameter λ that, differently for the other parameters, does not have a clear economic meaning; higher values of this parameter increase the speed of the learning process.

Figure 2.1 *left* and *right* shows average effort strategies learnt by workers when the share of permanent contracts is respectively 20 and 40 percent.¹⁰ We repeat a complete run of the simulation (500 periods) for the three values of λ and plot average strategies at the end of the simulations; effort is the x -axis and the probability of playing each effort value on the y -axis. The other values of the parameters are the ones in the baseline case,

<http://ec.europa.eu/eurostat/web/income-and-living-conditions/overview>

⁹When the share of permanent contracts is 0.5, 0.6, 0.7, 0.8, 0.9, 1 the standard deviation is respectively 0.84, 0.67, 0.53, 0.30, 0.13, 0.

¹⁰Similar results hold for all the other values of the share of permanent contracts; we omit them to save space.

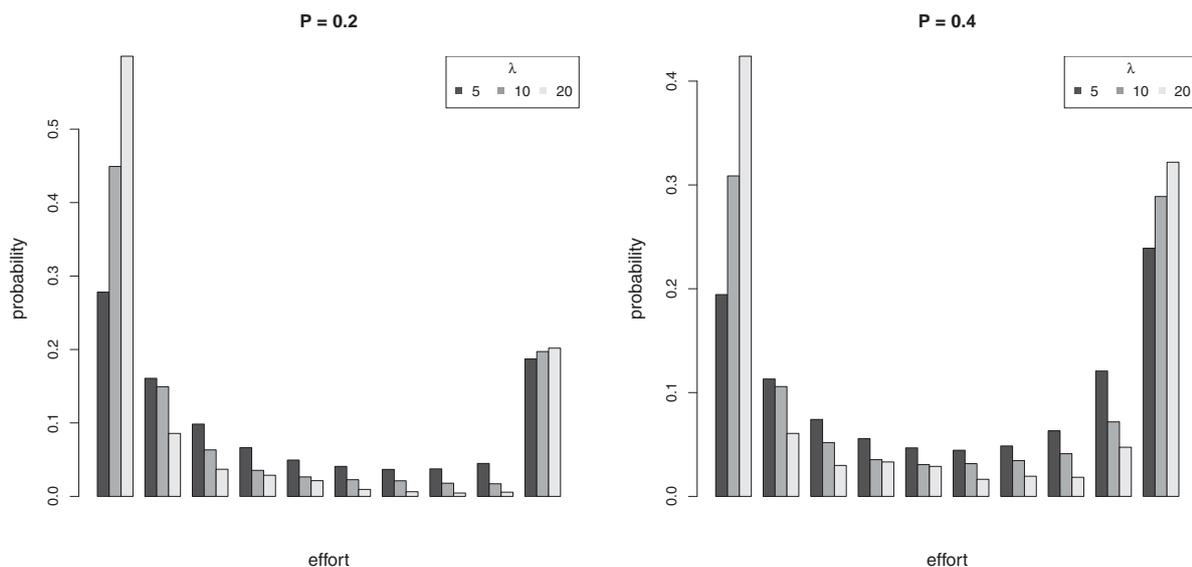


Figure 2.1: Temporary workers' average strategy for $P = 0.2$ (*left*) and $P = 0.4$ (*right*) for three different values of the parameter $\lambda = \{5, 10, 20\}$ determining the speed of the reinforcement learning process.

described in Table 2.1.

As shown in both plots, similar results are obtained with the three different values of the learning parameter and when $P = 0.4$ maximum (minimum) effort is played with higher (lower) probability compared to case $P = 0.2$. The main difference between the effort distributions with the different values of λ is in the probability of playing the minimum effort value $e = 1$. For low values of the learning process, when $\lambda = 5$, minimum effort is played with lower probability compared to the other two cases. This is an expected result as $\lambda = 5$ corresponds to a slower learning process and effort values are initially played with equal probability.

Table 2.1 summarizes the results of robustness checks for different values of the parameter λ with respect to average effort.

	Share of permanent contracts P									
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$\lambda = 5$	3.72	4.48	5.13	5.69	6.18	6.51	6.58	6.32	5.21	2.61
$\lambda = 10$	2.93	3.77	4.53	5.26	5.96	6.48	6.80	6.87	5.93	1.62
$\lambda = 20$	2.31	3.31	4.25	4.94	5.50	6.26	6.42	5.89	5.27	2.60

Table 2.1: Average effort for all values of the share of permanent contracts $P = \{0.1, 0.2, \dots, 1\}$ with three values of the learning parameter $\lambda = \{5, 10, 20\}$.

Each value corresponds to average effort at the end of simulation where the share of available permanent contracts P is indicated in the columns. Simulation with different values of λ lead to similar average effort values and, comparing values in the three rows, one can see that the inverted-U-shape result linking the share of permanent contracts P and average effort is robust across the simulations.

Figure 2.2 test the robustness of results with respect to the parameter α , the cost of one unit of effort. In the baseline case effort costs are drawn from a uniform distribution $U \sim [0.05, 0.15]$, here we instead simulate the model for three different values of $\alpha = \{0.05, 0.1, 0.15\}$.

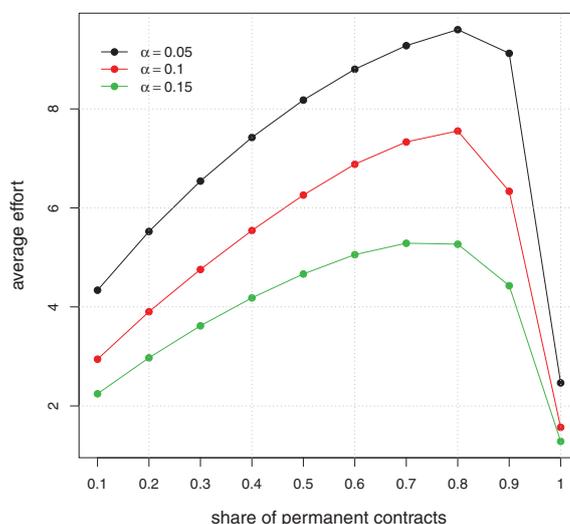


Figure 2.2: Share of permanent contracts P and average effort for three values of the effort cost parameter $\alpha = \{0.05, 0.10, 0.15\}$.

The figure summarizes average effort, at the end of a simulation, for all values of the share of permanent contracts, for three values of the effort cost parameters. The inverted-U-shape relationship between the share of permanent contracts and average effort still clearly appears for the three values of the parameter α and, in the three cases, 80% is the optimal share. Moreover, as expected, higher effort cost corresponds to lower average effort values.

2.5 Concluding remarks

This paper is a contribution to the recent and growing literature trying to assess the effect of temporary employment on labor market outcomes. We focus on behavioural aspects of

this phenomenon and, in particular, the goal of this paper is to assess whether, and under which conditions, temporary employment induces an increase in workers' willingness to exert high effort and, as a consequence, which is the effect on productivity. We therefore take a *micro* approach based on individuals' behaviour, focusing on effort as the *only* channel through which temporary employment affects productivity. Nevertheless, we are well aware of the presence of other potential channels such as firms' willingness to innovate or invest in human capital.

We implement an agent-based model where workers and firms interact in a dual labor market, with temporary and permanent contracts. The agent-based model allows us to take into account behavioural aspects previously neglected in the literature; to the best of our knowledge this is the first paper applying this methodology to study dual labor markets.

In the model, within firms workers with temporary contracts compete for permanent positions and they face the following *trade-off*: exerting high effort is a costly investment but it increases the chances of getting a contract upgrade. Using reinforcement learning, workers revise their effort strategies on the basis of past experience. The main result is that temporary workers' optimal effort depends on the share of available permanent contracts; the relationship between the share of permanent contracts and effort, and consequently labor productivity, has an inverted-U-shape. As the share of available permanent contracts increases, workers optimally increase effort but, when too many permanent positions are offered, workers have an incentive to shirk as they are too confident of being upgraded.

Using a different methodology, Hirsch and Mueller (2012) investigate the effect of temporary agency work on firm productivity, allowing for a flexible relationship between the two variables. Using a large panel of German plants the authors find a non-linear hump-shaped relationship between productivity and temporary agency work use, with the maximum positive effect on productivity when hiring about 11 percent of temporary external workers.

Our results should be taken as suggestive rather than conclusive, and the model is very simple and therefore could be improved in several dimensions. For example, firms cannot choose or adjust the share of permanent contracts, as it is an exogenous parameter, but it could instead be a firm's choice variable. Therefore, one could think of a model in which firms choose the share of permanent contracts that maximizes profits, taking into account workers' effort responses.

Moreover, temporary workers do not consider the probability of finding a new job once unemployed. If, say, there are many unemployed and it is hard to find a new job workers might be willing to exert higher effort not to bear the risk of remaining unemployed for

a long time; in other words market tightness is not taken into account.

For OECD countries the transition probability that a temporary contract is converted into permanent is relatively small, never larger than 50% and as low as 12%-13% in Portugal and France, our model instead predicts an optimal share greater than 50%. Therefore, the model suggests that converting too few contracts into permanent may be providing workers incentives to exert low effort and, consequently, this may be one of causes of low productivity values, recently experiences in some European countries. Therefore, while the European debate on labor market reforms is mainly focused on net employment effects, our results suggest that a proper regulation design requires considering factors such as the effects on workers' behaviour.

Tesi di dottorato "Essays on Labour Economics and Education"
di FANO SHIRA

discussa presso Università Commerciale Luigi Bocconi-Milano nell'anno 2016

La tesi è tutelata dalla normativa sul diritto d'autore(Legge 22 aprile 1941, n.633 e successive integrazioni e modifiche).

Sono comunque fatti salvi i diritti dell'università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte.

Appendix

2.A Empirical evidence

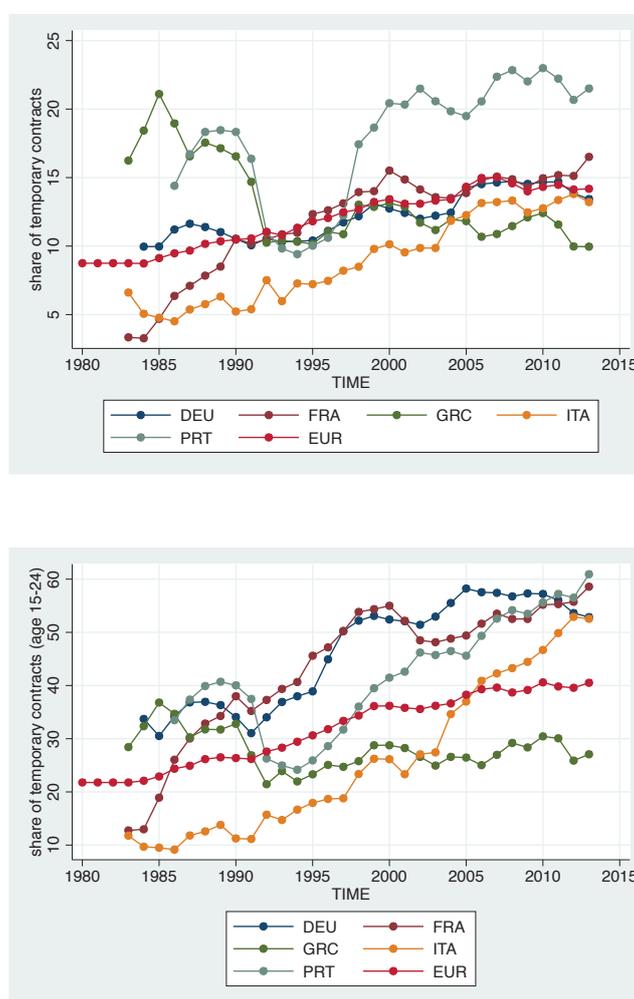


Figure 2.A.1: Share of temporary contract (*top*) share of temporary contracts for the age class 15–24 (*bottom*); both excluding Spain.

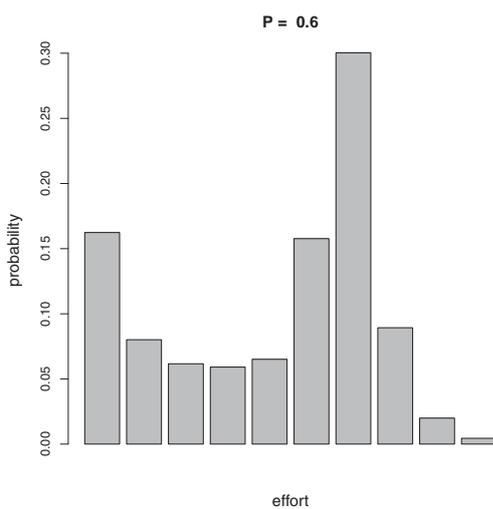
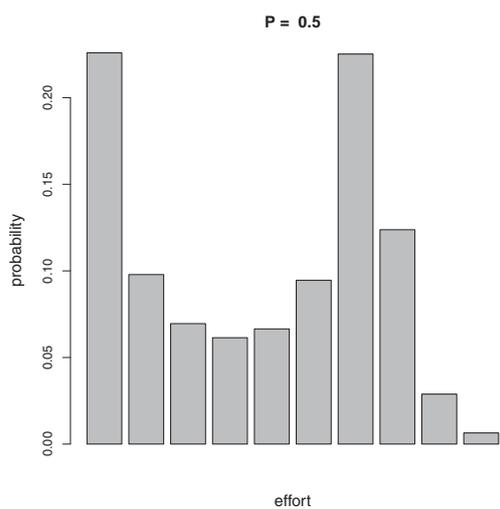
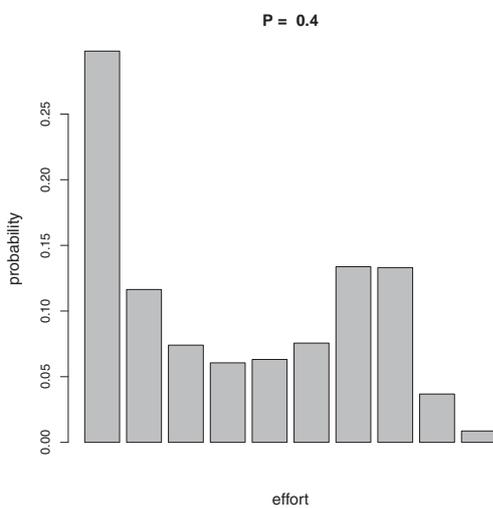
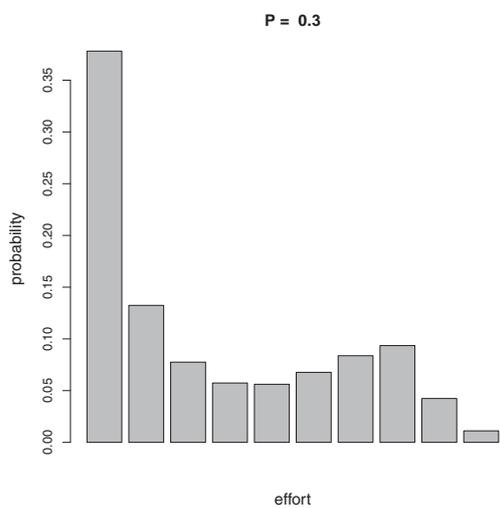
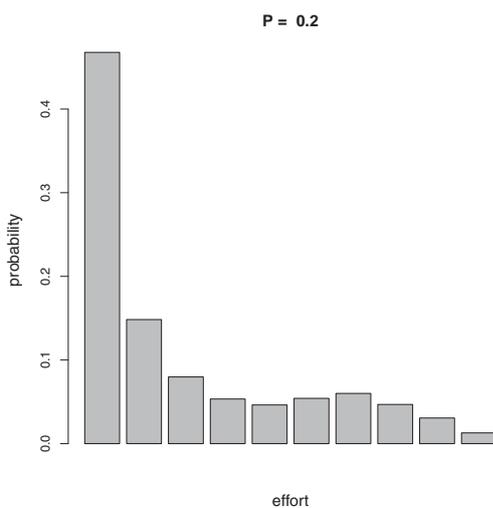
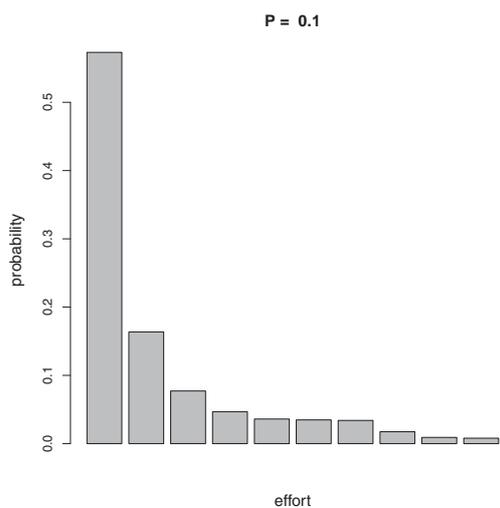
2.B Additional simulation results

Share P	Probability of playing effort k									
	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
0.1	0.51	0.18	0.07	0.04	0.03	0.02	0.02	0.01	0.01	0.11
0.2	0.45	0.15	0.06	0.04	0.03	0.02	0.02	0.02	0.02	0.20
0.3	0.38	0.13	0.06	0.04	0.03	0.03	0.03	0.03	0.03	0.26
0.4	0.31	0.11	0.05	0.04	0.03	0.03	0.03	0.04	0.07	0.29
0.5	0.24	0.09	0.05	0.03	0.03	0.04	0.04	0.06	0.14	0.30
0.6	0.18	0.07	0.04	0.03	0.03	0.04	0.05	0.11	0.22	0.23
0.7	0.11	0.05	0.04	0.03	0.04	0.04	0.10	0.22	0.23	0.14
0.8	0.06	0.03	0.03	0.03	0.03	0.08	0.28	0.26	0.13	0.06
0.9	0.01	0.01	0.01	0.07	0.28	0.33	0.17	0.07	0.03	0.01
1	0.64	0.22	0.08	0.03	0.02	0.01	0.00	0.00	0.00	0.00

Table 2.B.1: Average strategy learnt by all workers at the end of a simulation for different values of the share of $P = \{0.1, 0.2, \dots, 1\}$.

P	Median	Mean	sd
0.1	2.50	2.70	0.85
0.2	3.20	3.11	0.81
0.3	3.45	3.57	0.98
0.4	3.55	3.60	0.69
0.5	3.55	3.48	0.84
0.6	3.20	3.20	0.67
0.7	2.70	2.69	0.53
0.8	2.30	2.21	0.30
0.9	1.50	1.51	0.13
1	1	1	0

Table 2.B.2: Average productivity for different values of the share of permanent contracts $P = \{0.1, 0.2, \dots, 1\}$, baseline case.



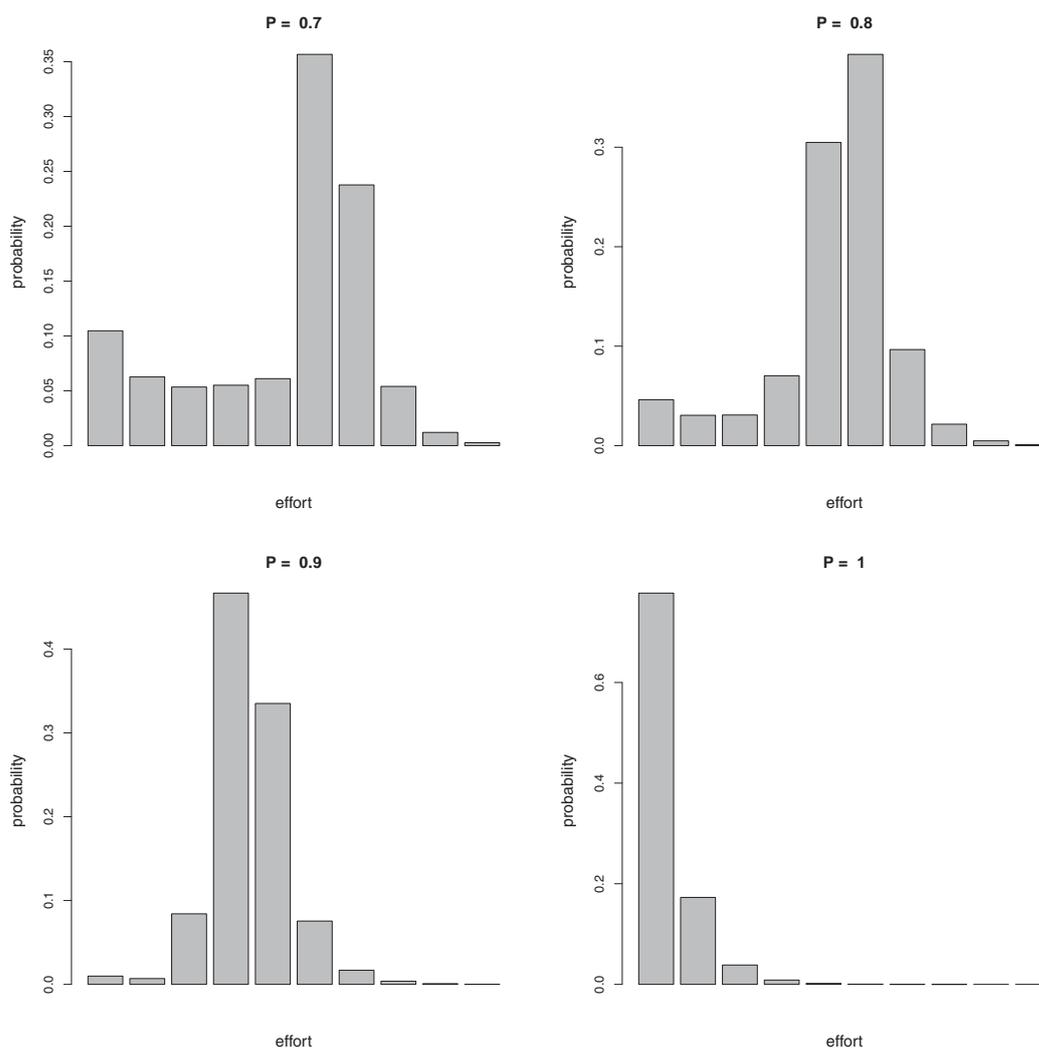


Figure 2.B.1: Workers' average strategy at the end of a simulation for all possible values of the share of permanent contracts in $P \in \{0.1, 0.2, \dots, 1\}$ when $\alpha = 0.15$, the other parameters are those in the baseline case.

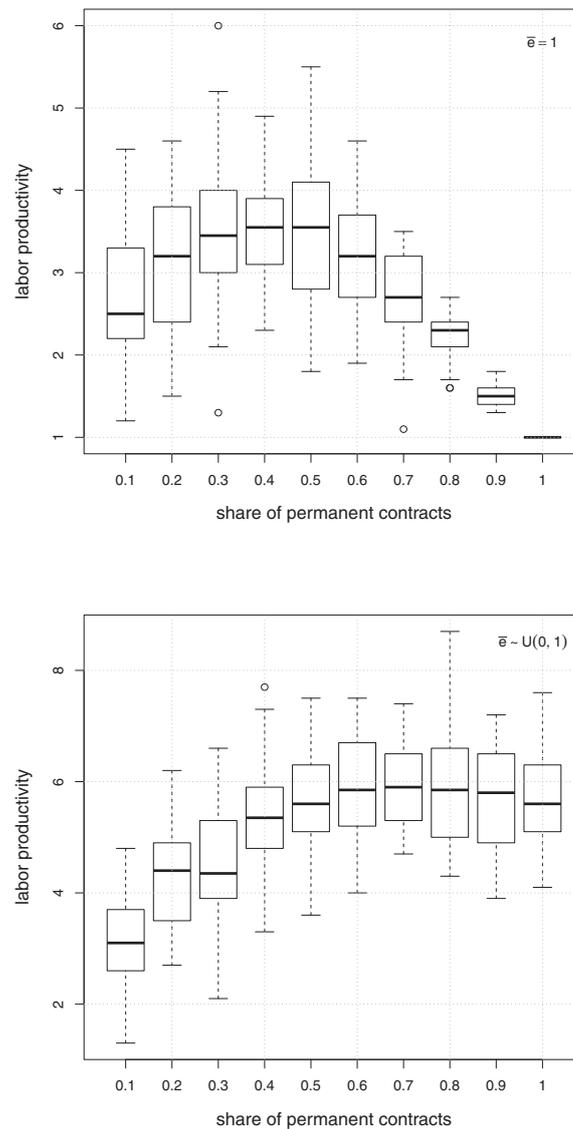


Figure 2.B.2: Share of permanent contracts and labor productivity when permanent workers supply minimum (*left*) and 0.5 (*right*) effort; parameters of the baseline case.

Tesi di dottorato "Essays on Labour Economics and Education"
di FANO SHIRA

discussa presso Università Commerciale Luigi Bocconi-Milano nell'anno 2016

La tesi è tutelata dalla normativa sul diritto d'autore(Legge 22 aprile 1941, n.633 e successive integrazioni e modifiche).

Sono comunque fatti salvi i diritti dell'università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte.

Chapter 3

After the crisis, are workers allocated to more productive firms? Evidence from the CompNet dataset

Abstract

This paper investigates the effect of the financial crisis on the allocation of workers across firms, using an index of allocative efficiency that links firms size and productivity called the OP-gap.

We use a newly constructed database by the members of the Competitiveness Research Network (CompNet) at the ECB that provides information on the distribution of firms across different dimensions related to competitiveness, such as size and productivity. We focus on Belgium and Spain, for which we have a representative sample of the population of firms and data from 1995 to 2011.

The main finding is that the financial crisis had a positive and significant effect on allocative efficiency only in Spain and not in Belgium and, after the crisis, Spanish firms are able to better allocate workers to more to productive firms. Distinguishing between tradable and non-tradable sectors, after the financial crisis allocative efficiency in Spain decreases in tradable sectors and increases in non-tradable ones.

Keywords: Allocative efficiency, productivity analysis, financial crisis.

3.1 Introduction

Both theoretical and empirical literature have shown that there are large and persistent differences in productivity across countries, and an increasing number of papers has tried to identify the drivers of this heterogeneity.

At the same time, in the last decade a parallel strand of literature has emerged showing that there is large heterogeneity in productivity also among different firms within the same sectors and countries, see Bartelsman et al. (2004).

In recent years few papers have combined these two strands of research by trying to explain how cross country differences in productivity may be related to differences in the within-industry productivity dispersion across firms. The idea is that the observed heterogeneity in firm-level productivity may be due to a misallocation of resources across firms and, in turn, this may have a negative effect on aggregate productivity levels. In general terms, allocative efficiency refers to a situation where available resources are put to their best use. However, which measures of firm level heterogeneity are more instructive to detect possible mis-allocations of resources is still an open question.

In other words, cross-country productivity differences partially can be explained by differences in allocative efficiency and aggregate productivity in a country may be lagging partly because available inputs are not allocated efficiently across firms within an industry.

This finding provides a potentially new channel for boosting aggregate productivity, through the reallocation of resources away from poorly productive firms towards the most productive ones. This hypothesis of misallocation of resources as a source of low productivity is not new (see, e.g., the handbook paper by Banerjee and Duflo (2005) for a review) but the development of firm-level databases in a variety of countries now permits exploring this hypothesis more directly.

In this paper, we explore the misallocation hypothesis and show how allocative efficiency varies across countries, more interestingly, over different periods of time. The main contribution of this paper is to study the effect of the financial crisis on the allocation of workers across firms, by looking at changes over time of an index of allocative efficiency based on the within-industry covariance between firm size and productivity.

To quantify the within-industry covariance between size and productivity, that is our proxy of allocative efficiency, we use an established empirical decomposition of the level of industry productivity as proposed by Olley and Pakes (1996), (henceforth OP). The OP decomposition splits the weighted average of firm-level productivity into (i) an unweighed firm-level average and (ii) a covariance term, called OP-gap, that is a summary measure of the within-industry cross sectional covariance between size and productivity. The

covariance term could be interpreted as the degree to which resources are allocated efficiently across firms within the same industry. A low covariance indicates that aggregate productivity can improve by reallocating resources towards the most productive firms.

Although we are still far from understanding why allocative efficiency varies across sectors, and countries, there is a growing number of papers relating those observed differences to sector specific regulations of the labour, the product or credit market (see for example Restuccia and Rogerson (2008), Arnold et al. 2011, Akerberg et al. (2012), Aghion et al. (2007) and Martin and Scarpetta (2011)).

Our data source for the econometric analysis is a newly constructed sectoral database build by the members of the Competitiveness Research Network (CompNet) at the ECB that provides information on the distribution of firms across different dimensions related to competitiveness, such as size and productivity. The construction and validation of the database is still in progress and, in what follows, we focus on two countries, Belgium and Spain, for which we have a representative sample of the population of firms and data from 1995 to 2011.

In their seminal contribution, Olley and Pakes (1996) found that the covariance term (using a decomposition of industry Total Factor Productivity, TFP) increased substantially in the US telecommunications equipment industry following the deregulation of the sector in the early 1980s. They argued that this was because the deregulation permitted outputs and inputs to be reallocated more readily from less productive to more productive US firms. Melitz (2003) using firm-level data provides evidence that there is a positive correlation between the distributions of productivity and size that is to say, more productive firms tend to be larger than less productive ones. However, the cross-country data analysed suggest that there is considerable variation in the strength of the link between productivity and size across countries and industries and over time. More recently, Bartelsman et al. (2013a) confirm the existence of a substantial and systematic cross-country variation in the within-industry covariance between size and productivity. They develop and calibrate a model to show how differences in size-productivity covariances imply substantial differences in aggregate performance.

The rest of the paper is organized as follows. Section 3.2 describes the aim, methodology, characteristics and critical points related to the construction of the CompNet database. Section 3.3 describes labor productivity and section 3.4 defines the allocative efficiency index, and describes its time series and sector heterogeneity. Section 3.5 analyses the effect of the financial crisis on allocative efficiency and section 3.6 relates allocative efficiency to sector shares. Section 3.7 concludes and suggests further research.

3.2 The CompNet Database

The aim of this section is simply to describe procedures followed by the CompNet project to build and validate a dataset of comparable cross-country indicators that can be systematically used for policy purposes. In this paper, we will then exploit data from the newly constructed dataset for the econometric analysis.

Cross-country analysis are crucial tools for policy making as they allow to benchmark and look for best practices in peer countries and they have therefore widely and increasingly been used by international organizations such as the OECD and World Bank. Cross-country comparisons allow for example to investigate the impact of similar shocks on different economies characterized by specific economic institutions and market structures.

The economic literature has shown that firm-level data is crucial information for understanding the drivers of competitiveness, because the aggregate performance of the economy strongly depends on characteristics and decisions at the firm level and shocks might have a different impact depending on the underlying distribution of firms in the economy.

For these reasons, the Competitiveness Research Network of the EU System of Central Banks (CompNet), that was established with the aim of analysing the determinants of country competitiveness, has put micro-data and sector-level analysis at the center of its interests.

However, in practice cross-country analysis is hindered by two main reasons. First, the existing indicators based on firm level data are often not comparable because they refer to different periods and moreover, they have been collected using different procedures and definitions of variables. Second, information at the firm level is usually highly confidential. For those reasons, the commercial databases compiled by Bureau van Dijk (Amadeus is the European version of it) with information from the firm registries have been widely used. However, the drawback of these databases is that some variables, like employment, are not compulsory in all countries so actual firm coverage is drastically reduced.

To set-up a new research infrastructure able to deliver cross-country firm-based indicators, CompNet has followed the approach known as “distributed micro-data analysis”, that was previously followed by the World Bank and the OECD in projects on firm dynamics, see Bartelsman et al. (2009). To ensure comparability across countries a common set of codes were prepared at the Central Bank and the same codes were executed by the National Central Banks by different country teams on their firm-level datasets, taking advantage of the link between the ECB and the National Central Banks (NCB).

The common codes allowed to aggregate the firm-level data at the sector-level, thus preserving firm-level confidentiality. The CompNet dataset is therefore a sectoral dataset and the output of the exercise is a database featuring 11 EU countries and 58 NACE Rev.2 industries over the period 1995-2011, with comparable information on productivity performance and dynamics of underlying heterogeneous firms. The countries included are: Belgium, Czech Republic, Germany, Estonia, France, Hungary, Italy, Poland, Spain, Slovakia and Slovenia and together they represent approximately two-thirds of European Union's GDP. The Italian National Statistical Institute (ISTAT) and the EFIGE team were also involved.¹

Special attention and effort was put to ensure that the protocol treated the data in the same way on a number of crucial issues including the classification of sectors, the use of deflators, the treatment of outliers and the definition and computation of variables.

The advantage of CompNet's research infrastructure with respect to existing information included in aggregate statistics is that it enables to keep much of the richness of firm-level data in terms of full distribution of variables or joint correlations computed at the firm-level. With respect to existing work based of firm-level information, but reported in aggregated fashion (e.g. Eurostat), the CompNet data is able to exploit the information content coming not only from averages, but also from the distribution of firms across several dimension, e.g. productivity and size.

Several rounds of data collection have taken place, errors were corrected and more variables were added to the dataset. In this paper, we will use available data from the 3rd round of data collection. Out of the 11 listed countries, 5 have reasonably representative samples in terms of size distribution and sector whereas 6 have an over-representation of large and manufacturing firms (Slovakia, Poland, Italy, Hungary, Germany and Czech Republic). To partially mitigate the bias towards large and manufacturing firms of some country samples, a common set of sector weights were created to compute country aggregates. The sector weights are computed as the average value added share of each industry across the 11 countries.

The National teams were also in charge of providing detailed information on the meta-data, that is, the description of their databases, their sources and existing thresholds, regarding for example confidentiality issues. A validation process of the raw data has also been done by comparing data collected by the CompNet teams with the *Structural*

¹EFIGE (European Firms in a Global Economy) is a project that has recently developed a firm level database, including 7 European countries, supported by the Directorate General Research of the European Commission through its 7th Framework Programme and coordinated by the European think thank Bruegel.

Business Statistics of the Eurostat. The exercise consists in computing the correlation over time between variables such as total turnover, value added, employment and labour cost aggregated at the country/industry/year level from CompNet with those from official statistics provided by Eurostat. In general the values of the correlation between the CompNet and the Eurostat data are very high, with most values above 0.8 and many over 0.9. Overall, the validation exercises showed that the sector-level aggregated variables computed from the CompNet data show characteristics and dynamics consistent with those of other aggregated sources.

3.3 Labor productivity

In the paper we look at allocative efficiency in two countries: Belgium and Spain. Since, in what follow we will proxy allocative efficiency with a component of labor productivity, we start by simply comparing the value of labor productivity in the two countries, using data collected in the CompNet dataset. Labor productivity is computed as real value added divided by employment, expressed in thousands of euros per employee. Alternative productivity measures are also computed as the ratio between turnover and employment.

Figure 3.1 depicts the values of labor productivity in Belgium and Spain, grouped between tradable and non tradable sectors, that roughly correspond to manufacturing and services. More precisely, tradable sectors include manufacturing (NACE Section C) with the exclusion of “Manufacture of coke and redefined petroleum products” and non tradable sectors include construction (NACE Section F), accommodation and food service activities (NACE Section I), information and communication (NACE Section J), professional, scientific and technical activities (NACE Section M) and administrative and support service activities (NACE Section N). Sector weights used in the aggregation procedure are those described in the appendix, but computed separately for tradables and non-tradables so they sum up to one for each of the aggregates. Then, an unweighted average over the years 2003-2007 is computed.

Note that the scale on the y -axis is different in the two plots and, as expected, labor productivity exhibits higher values in Belgium when compared to Spain. Although looking at the country mean labour productivity, CompNet firm-level data replicate well known rankings calculated at the macro (aggregate) level across countries, cross-country comparisons of labour productivity levels have to be done with lots of caution. In fact, labour productivity measures are not expressed in terms of Purchasing Parity Power (PPP), but in thousands of euros. This can be driving some of the difference in productivity levels and, expressing productivity in comparable PPP units for Spain would result in Spanish

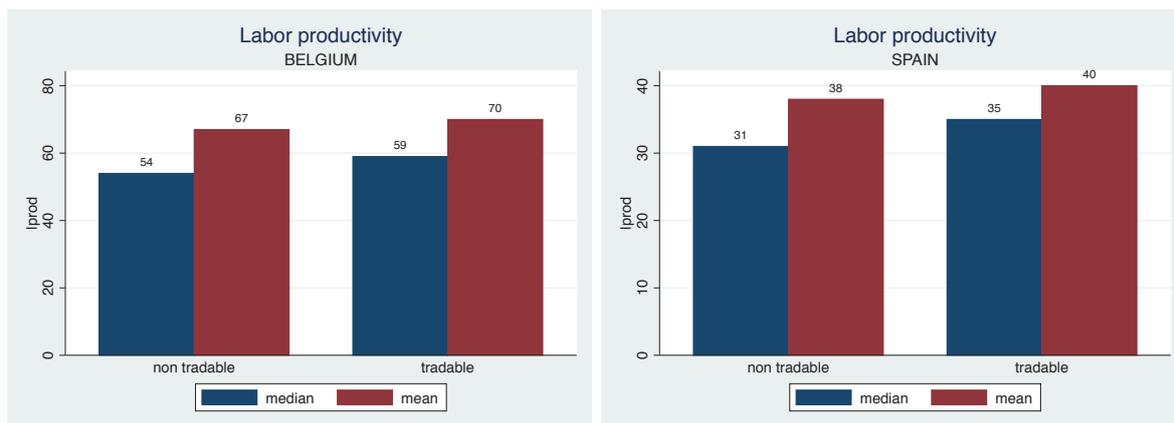


Figure 3.1: Labor productivity in Belgium (right) and Spain (left) in tradable and non-tradable sectors.

firms having a slightly higher productivity than reported, but still significantly below the level of Belgium.

In both countries, the mean is above the median, both in tradable and non-tradable sectors. This reflects the skewness of the distributions that are characterised by a large group of low productivity, small firms and few champions, and therefore exhibit long right tails.

Moreover, both in Belgium and Spain, labor productivity values are higher in tradable compared to non-tradable sectors.

3.4 Measuring allocative efficiency

In the literature there are two definitions of allocative efficiency: the first is static and the second is dynamic. In what follows, we focus only on the first concept. Static allocative efficiency is a measure of how, in a given moment in time, firms with higher than average productivity have a larger size in terms of employment. Olley and Pakes (1996) were the first to develop this index, based on the covariance between size and productivity at the industry level.

According to the index aggregate productivity can be decomposed as follows:

$$y_{st} = \sum_{i \in s} \theta_{it} \omega_{it} = \bar{\omega}_{st} + \underbrace{\sum_{i \in s} (\theta_{it} - \bar{\theta}_{st})(\omega_{it} - \bar{\omega}_{st})}_{\text{OP Gap}} \quad (3.1)$$

where y_{st} is the weighted average productivity of sector s at time t , θ_{it} and ω_{it} are respectively size and productivity of firm i at time t and $\bar{\theta}_{st}$ and $\bar{\omega}_{st}$ are respectively the

unweighted average size and productivity of sector s at time t . Hence allocative efficiency is proxied by the covariance between the relative size of a firm and its relative productivity. Higher is the covariance term, better are resources allocated.

If resources were allocated randomly across firms in the industry the covariance measure in the right-hand side of equation (3.1) would be zero, and aggregate and average productivity would coincide. The larger the covariance, the more efficiently are resources allocated within the sector and the higher the contribution of the (efficient) allocation of resources to the sector productivity, vis-à-vis the unweighted average productivity of the firms operating in the sector.

The static concept of allocative efficiency provides a snapshot of how resources are allocated at a certain moment in time. Hence, the covariance term could be interpreted as the degree to which resources are allocated efficiently across firms within the same industry. A low covariance indicates that aggregate productivity can improve by reallocating resources towards the most productive firms. See Bartelsman et al. (2009) for a thorough analysis of different measures of allocative efficiency.

Figure 2 depicts the OP-gap for Belgium and Spain, grouped in tradable and non tradable sectors, the figure therefore shows the last term on the right-hand side of equation (3.1).

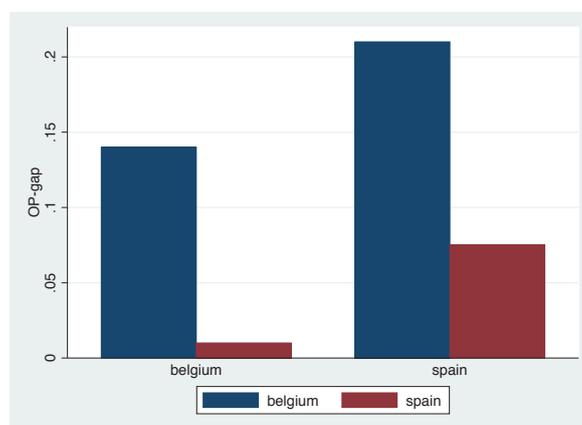


Figure 3.1: OP-gap in Belgium (right) and Spain (left) in tradable and non-tradable sectors.

The covariance is first computed at the 2-digit industry level and then aggregated over all tradable and non-tradable sectors in each country. Sector weights are those described in the appendix, but computed separately for tradables and non-tradables so they sum up to one for each of the aggregates. Then, an unweighted average over the pre-crisis years 2003-2007 is computed.

Notice first that numbers we obtain are generally quite low, although consistent with respect to those shown in other papers such as Bartelsman et al. (2009). The covariance goes up to 0.2 in Spain, which means that sector labour productivity is up to 20% higher than it would be with randomly allocated labour. Hence, from an accounting perspective, it is clear that the contribution of the unweighted average productivity of the firms operating in the sector to aggregate productivity is larger than that of the covariance term. Or in other words, the contribution of the allocation of resources to overall productivity is lower than the contribution of the average productivity of the existing firms.

However, recent theoretical (Acemoglu et al. (2013)) and empirical (Andrews and Cingano (2012)) contributions, show that the incentive for firms to increase their own productivity goes up when resources flow easily to the best firms. This means that the direct contribution from better allocation is boosted through its indirect effect on within-firm productivity growth.

Secondly, there are large within-country differences in terms of allocative efficiency in tradable vs. non-tradable sectors. This large difference in terms of allocative efficiency between tradables and non-tradables is not surprising, and it has also been found in other works like Arnold et al. (2011). They might reflect the fact that regulatory reforms in non-tradables have been more hesitant, above all in mature European countries, and that these sectors are sheltered from competition.

Cross-country differences in allocative efficiency are also remarkable. Spain feature a larger covariance between size and productivity, that is, a higher allocative efficiency both in tradables and non-tradables, compared to Belgium. This means that the most efficient firms may be large but recall that they are not as productive as the best firms in Belgium, see Figure 3.1. It is beyond the scope of this paper to discuss what policies could be put in place to increase allocative efficiency, although an increasing number of papers on the topic are pointing to labour and product market regulations as well as financial development or a fair and predictable legal system, see Arnold et al. (2008) and Haltiwanger (2011).

3.4.1 Time series

The following two charts show the time series of the OP-gap respectively for Belgium (left) and Spain (right) in the time period 1995-2011.

The Op-gap is plotted for tradable and non-tradables and the vertical bar in 2008 is meant to point out the beginning of the financial crisis. As previously discussed, the covariance term is first computed at the two digit industry level and then aggregated using

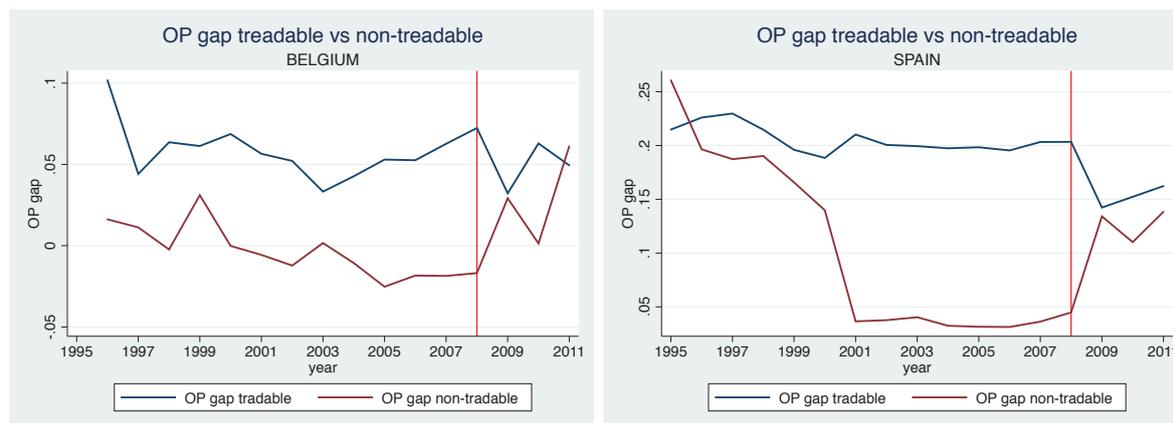


Figure 3.2: Time series of the OP-gap in Belgium and Spain in 1995-2011.

real value added as a weighting variable; weights are computed separately for tradables and non-tradables.

Note that the scale on the vertical axis is different across charts. For both countries, the value of the OP-gap for tradables is higher than the value of the index for non-tradables. For tradables, the value of the index is higher in Spain than in Belgium, approximately 0.2 and 0.05 respectively. For both countries, there is a drop in the index in 2008 for tradables of approximately 0.05. An interesting pattern emerges for Spain. Between 1995 and 2005 the two time series follow a similar pattern but they diverge between 2000 and 2001 when the value for non-tradables drops to approximately 0.05.² After this period, the OP-gap is basically stable until the financial crisis. In 2008, the two time series converge as the index decreases for tradables and increases for non tradables. For non-tradables the index jumps from approximately 0.05 to 0.15 in one year. This means that before the crisis, the allocation of resources to more productive firms increased productivity by 5%, whereas, after the crisis, the productivity gain due to the allocation of resources to more productive firms is 15%.

3.4.2 Sector analysis

We take a sector approach and check if some of the sectors are driving the aggregate pattern. The following four graphs plot the OP-gap at the two digit industry level in Belgium (top) and Spain (bottom). In each of the four graphs we split the available years

²This corresponds to the period of the dot-com bubble, when a very high number of internet-based companies were founded. Their stock prices increased rapidly since investors were confident to gain future profits but this created a bubble that collapsed in 1999-2001. Many companies failed and others lost big shares of the market.

in two groups: pre and post-crisis. The first group is an unweighted average of the years 2003-2007, the second of the period 2008-2011.

Again, comparing tradables and non tradables, for both Belgium and Spain, tradable sectors are able to better allocate workers across firms. This difference in terms of allocative efficiency has been found, for example, also in recent work by the European Commission (2013). Moreover, in most sectors the value of the OP gap is higher in Spain than in Belgium.

Both pre and post-crisis in Spain in all of the tradable sectors the covariance term between size and productivity is positive, thus the allocation of resources contributes positively to aggregate productivity. In particular, after the crisis the sector manufacture of beverage is the one with the highest value while manufacture of leather and related products is the one with the lowest value within Spanish tradable sectors. For Belgium, the value of the OP-gap has negative values in three sectors: manufacture of tobacco products, manufacture of furniture and repair and installation of machinery and equipment. Regarding non-tradables, more sectors have negative values of the covariance term be-

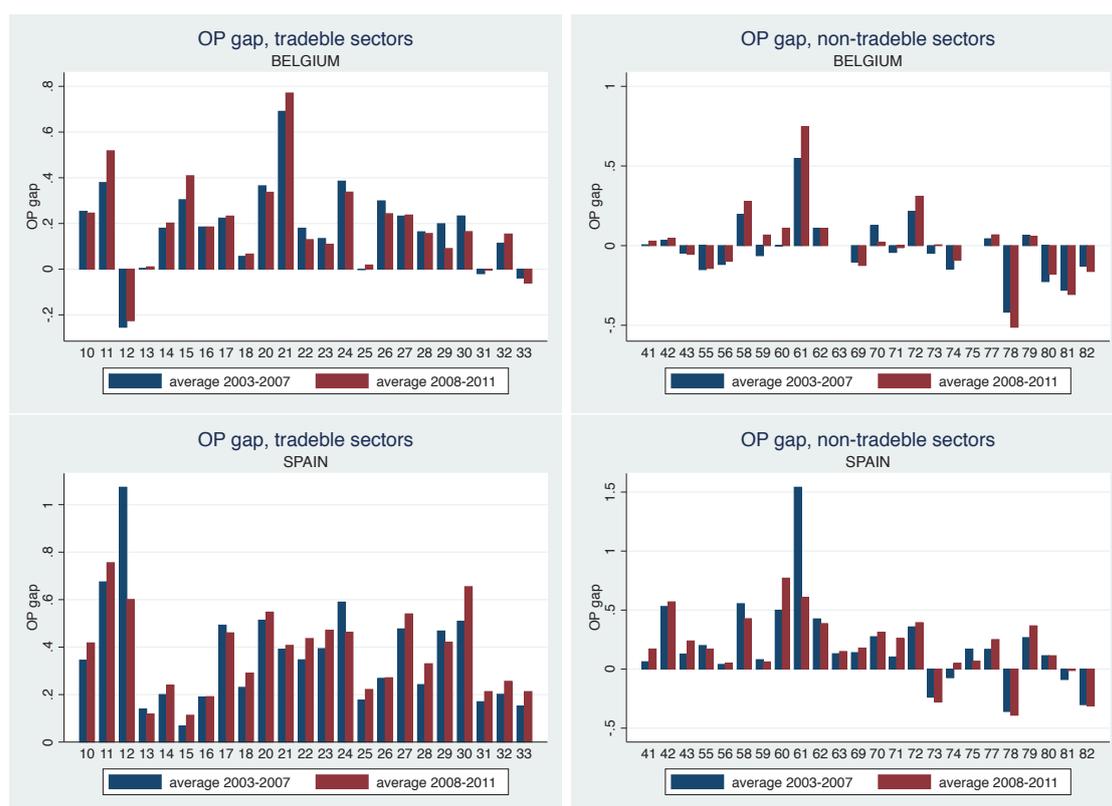


Figure 3.3: Allocative efficiency by sector, before/after the crisis.

tween size and productivity. In particular, in Spain the non-tradable sectors with negative values are: advertising and market reserve, employment activities and office administra-

tive office support and other business support activities.

Regarding the crisis, comparing pre and post bars the two countries seem to exhibit a different behaviour. In particular in Spain after the crisis the value of OP-gap increases in almost all sectors. The only three tradable sectors in which the OP-gap decreases after the crisis are 12, 17, 24 and 29 respectively manufacture of tobacco, paper, basic metals and motor vehicles trailers and semi-trailers. In Belgium instead this is not the case in sectors 12 and 17, but many other sector exhibit a decrease in the index.

3.5 Did the crisis affect allocative efficiency?

In this section we perform an econometric analysis to study the impact of the financial crisis on the OP-gap, in Belgium and Spain. The aim of the following regressions is therefore simply to check weather the financial crisis had a significant impact on the allocation of resources across firms and, if so, in which direction in the two countries.

We start by estimating the following OLS regression model:

$$OPgap_{cst} = \beta_1 + \beta_2 crisis + b_s + crisis * b_s + a_c + c_t + \varepsilon_{sct} \quad (3.2)$$

The *OP-gap* is the dependent variable and *crisis* is the main explanatory variable. It is a dummy variables that takes value 0 in the years 2003-2007, previous to the crisis, and it is equal to 1 in the period 2008-2011. As a set of covariates we control for sector, country and year specific effects, including dummy variables. We also include all the interaction terms between the variable crisis and the sector dummies, to control for sector specific effects of the financial crisis.

The estimates of regression 3.2 are summarized in Table 3.1. The first two columns refer to Belgium while columns 3 and 4 refer to Spain. Columns 1 and 3 show that, the crisis dummy is positively and significantly correlated to the OP-gap in Belgium and not in Spain. Estimates thus confirm that in Spain after the financial crisis the allocative efficiency index increased, indicating a better allocation of workers across firms.

Controlling for the other covariates, in Spain after the crisis the value of the OP-gap increases by 0.41. Regarding the other covariates, for space matters we decided to display only coefficients for four manufacturing sectors, but all sector are included in the regression, as indicated in the table. Dummies are statistically significant for most sectors in Spain, while this is not the case for Belgium, the same holds for the crisis-sector interactions. Columns 2 and 4 are simple robustness test and the only difference with columns 1 and 3 is the inclusion of year dummies; the results from a qualitative and numerical point of view do not change.

Table 3.1: Effect of the crisis on the OP Gap in Belgium and Spain

	<i>Dependent variable: OP Gap</i>			
	Belgium	Belgium	Spain	Spain
Crisis	0.00000 (0.065)	0.022 (0.068)	0.416*** (0.088)	0.355*** (0.091)
sector10	0.243*** (0.052)	0.235*** (0.052)	0.848*** (0.060)	0.848*** (0.060)
sector11	0.363*** (0.052)	0.355*** (0.052)	1.114*** (0.060)	1.114*** (0.060)
sector12	-0.074 (0.052)	-0.082 (0.052)	1.186*** (0.060)	1.186*** (0.060)
sector13	0.003 (0.052)	-0.004 (0.052)	0.607*** (0.060)	0.607*** (0.060)
sectorX
crisis*sector10	0.003 (0.086)	0.011 (0.086)	-0.385*** (0.124)	-0.385*** (0.123)
crisis*sector11	0.155* (0.086)	0.163* (0.086)	-0.313** (0.124)	-0.313** (0.123)
crisis*sector12	-0.151* (0.086)	-0.144* (0.086)	-0.540*** (0.124)	-0.540*** (0.123)
crisis*sector13	0.007 (0.086)	0.015 (0.086)	-0.443*** (0.124)	-0.443*** (0.123)
crisis*sectorX
Constant	-0.00000 (0.044)	-0.005 (0.046)	-0.461*** (0.043)	-0.398*** (0.046)
Sector FE	YES	YES	YES	YES
Course*Sector FE	YES	YES	YES	YES
Year FE	NO	YES	NO	YES
Observations	857	857	986	986
R ²	0.851	0.854	0.804	0.812
Adjusted R ²	0.829	0.830	0.778	0.783
F Statistic	39.873***	35.451***	31.067***	28.320***

Note:

*p<0.1; **p<0.05; ***p<0.01

In Table 3.2 we perform similar regressions explicitly adding the tradable, non-tradable dimension and we estimate the following regression:

$$OPgap_{cst} = \beta_1 + \beta_2 crisis + b_s + crisis * b_s + \beta_3 tradable + tradable * crisis + a_c + c_t + \varepsilon_{sct} \quad (3.3)$$

where the variable *tradable* is a dummy that takes value 1 for tradable sectors and 0 for the non-tradable ones. We also add the interaction term *tradable * crisis* to take into account the potentially differential effect of the financial crisis for tradable and non-tradable sectors on the OP-gap.

As in the previous set of regressions the outcome is more interesting for Spain compared the Belgium. Controlling for the other covariates also in this specification, for Spain the crisis dummy has a positive and significant effect of the OP-gap, and also the magnitude of the coefficients is the the same as in the previous table. Estimates confirm that tradable sectors exhibit higher OP-gap values, and the magnitude of the coefficient is 0.65 and statistically significant.

The interaction term *crisis * tradable* is negative and statistically significant, the estimated effect on the OP-gap is -0.39 and *p*-value is 0.012. This means that the joint effect of the crisis in tradable sectors weakens the positive marginal effects of the two dimensions crisis and tradable. This is consistent with Figure 1 which shows a negative effect of the financial crisis for tradable sectors and the positive effect on non-tradable ones.

For Spain, we perform two further robustness checks. To control for heteroskedasticity, we use a robust regression and we finally perform quantile regression estimating the conditional median instead of mean; results are summarized in Table 3.3. The magnitude of the crisis estimate increases respectively to 0.49 and 0.47 and results qualitatively do not change.

3.6 Allocative efficiency and sector shares

Following the recent Product Market Review of the European Commission (2013), we are interested in the relationship between allocative efficiency and market share of sectors. The idea is that a bad allocation of resources in a small sector has a minor impact on the overall economy compared to a bad allocation of resources in a key sector. In the following charts the size of sectors is measured in terms of real value added (labeled weights). The following four plots show the correlation between real value added, on the x-axis and the OP-gap, measuring allocative efficiency.

	<i>Dependent variable: OP Gap</i>	
	Belgium	Spain
Crisis	0.022 (0.068)	0.355*** (0.091)
tradable	-0.017 (0.052)	0.647*** (0.060)
crisis*tradable	-0.044 (0.086)	-0.390*** (0.123)
sector10	0.252*** (0.040)	0.200*** (0.060)
sector11	0.372*** (0.040)	0.466*** (0.060)
sector12	-0.065 (0.040)	0.539*** (0.060)
sector13	0.012 (0.040)	-0.041 (0.060)
sectorX
crisis*sector10	0.055 (0.080)	0.005 (0.123)
crisis*sector11	0.207*** (0.080)	0.077 (0.123)
crisis*sector12	-0.100 (0.080)	-0.151 (0.123)
crisis*sector13	0.059 (0.080)	-0.054 (0.123)
crisis*sectorX
Constant	-0.005 (0.046)	-0.398*** (0.046)
Sector FE	YES	YES
Year FE	YES	YES
Course*Sector FE	YES	YES
Observations	857	986
R ²	0.854	0.812
Adjusted R ²	0.830	0.783
F Statistic	35.451*** (df = 121; 735)	28.320*** (df = 130; 855)

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 3.3: Financial crisis and OP Gap, robustness tests

	<i>Dependent variable: OP Gap</i>	
	<i>robust linear</i>	<i>quantile regression</i>
Crisis	0.491*** (0.047)	0.471*** (0.156)
sector10	0.894*** (0.032)	0.870*** (0.148)
sector11	1.161*** (0.032)	1.168*** (0.148)
sector12	1.189*** (0.032)	1.171*** (0.232)
sector13	0.654*** (0.032)	0.651*** (0.148)
sectorX
crisis*sector10	-0.458*** (0.066)	-0.436*** (0.157)
crisis*sector11	-0.388*** (0.066)	-0.385** (0.165)
crisis*sector12	-0.570*** (0.066)	-0.426 (0.469)
crisis*sector13	-0.518*** (0.066)	-0.491*** (0.156)
crisis*sectorX
Constant	-0.509*** (0.023)	-0.509*** (0.148)
Sector FE	YES	YES
Year FE	YES	YES
Course*Sector FE	YES	YES
Observations	986	986
Residual Std. Error	0.065 (df = 870)	
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

The two top charts refer to Belgium, the bottom ones to Spain. We split the data in pre-crisis (2003-2007) in the left charts and post-crisis in the right charts. Finally, in each of the four charts we distinguish between tradable and non-tradable sectors.

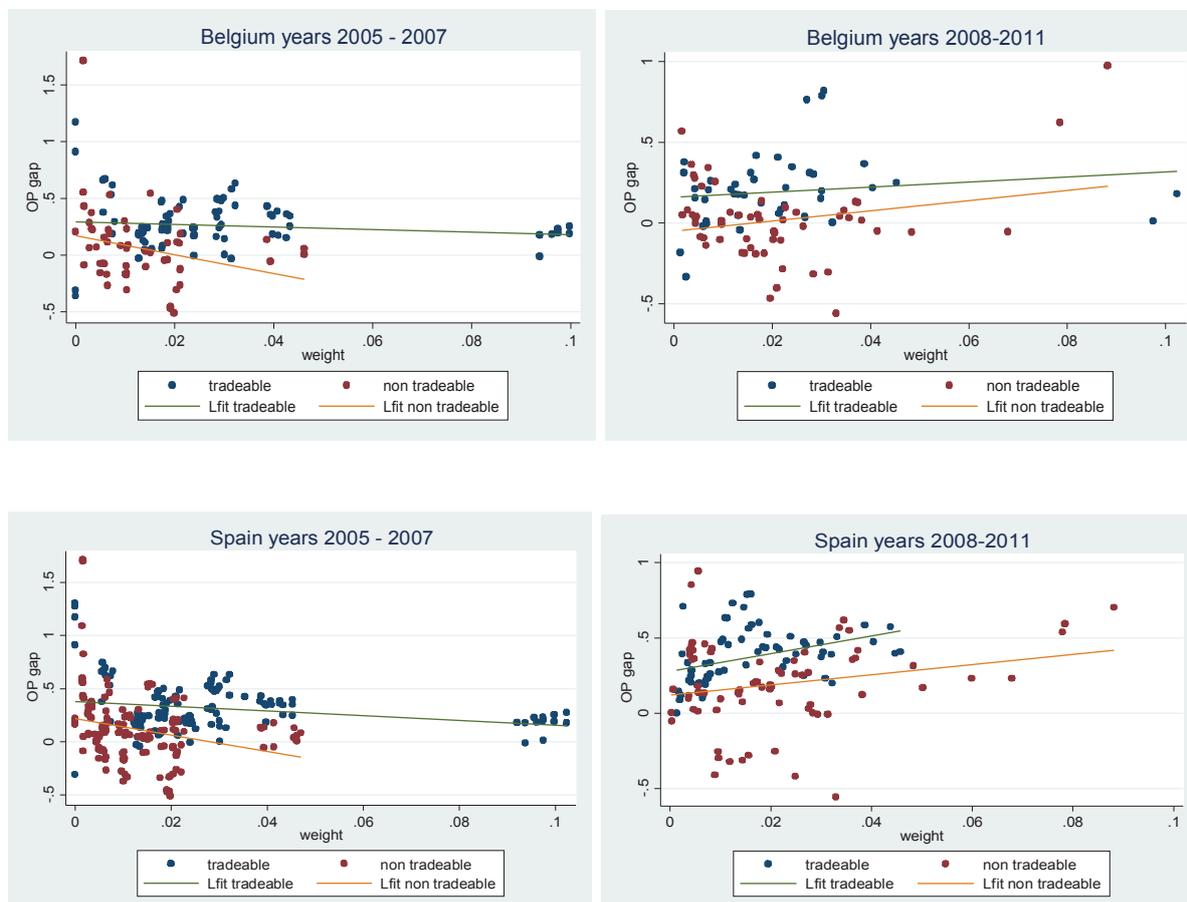


Figure 3.1: Allocative efficiency and market shares.

For both Belgium and Spain a similar pattern emerges. In the pre-crisis period the slope of the regression line is negative, meaning that sectors with greater real value added are the ones with lower values of allocative efficiency. In the post-crisis period the sign of the regression coefficients becomes positive both for Belgium and Spain and this effect is greater for Spain. This suggests that, after the crisis, bigger sectors in terms of real value added are also the ones better able to allocate workers across more productive firms.

3.7 Conclusion

This paper documents the construction of the CompNet database, an ambitious project of collection of harmonised firm-level based indicators, and presents some preliminary findings on the effect of the financial crisis on the OP-gap, a measure of allocative effi-

ciency. We focus on two countries: Belgium and Spain, for which we have a representative sample of firms in the period 1995-2011 and variables used are from the 3rd round of data collection.

The main results can be summarized as follows. The financial crisis had a positive and significant effect on allocative efficiency only in Spain, not in Belgium and, after the crisis, Spanish sectors further improve the allocation of resources to more productive firms. Distinguishing between tradable and non-tradable sectors, allocative efficiency decreased in tradable sectors and increased in non-tradables ones. Results are robust across different specifications.

Currently, in the dataset there is no information on the age of firms, so we cannot distinguish between incumbent and new firms. This is important when it comes to the analysis of resource reallocation given that one important channel of reallocation within sectors is the entry and exit of firms.

We are well aware that results should be taken as suggestive rather than conclusive and more analysis should be done, reducing for example potential biases due to omitted variables. The relationship between market shares and allocative efficiency should be further investigated. Nevertheless, for the case of Spain results suggest a so-called “cleansing effect” of the financial crisis. Previous to the crisis lots of inefficient firms were present in the market; an hypothesis is that the crisis harmed more the most inefficient firms that exited the market and boosting unemployment. Only the bigger and more efficient firms, able to better allocate resources, remained in the market; this can partly explain the increase in the index after the crisis for non-tradables.

Appendix

3.A Sector classification

5	Mining of coal and lignite	B	MINING AND QUARRING
6	Extraction of crude petroleum and natural gas	B	
7	Mining of metal ores	B	
8	Other mining and quarrying	B	
9	Mining support service activities	B	
10	Manufacture of food products	C	MANUFACTURING
11	Manufacture of beverages	C	
12	Manufacture of tobacco products	C	
13	Manufacture of textiles	C	
14	Manufacture of wearing apparel	C	
15	Manufacture of leather and related products	C	
16	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials	C	
17	Manufacture of paper and paper products	C	
18	Printing and reproduction of recorded media	C	
19	Manufacture of coke and refined petroleum prod	C	
20	Manufacture of chemicals and chemical products	C	
21	Manufacture of basic pharmaceutical products and pharmaceutical preparations	C	
22	Manufacture of rubber and plastic products		
23	Manufacture of other non-metallic mineral products	C	
24	Manufacture of basic metals	C	
25	Manufacture of fabricated metal products, except machinery and equipment	C	
26	Manufacture of computer, electronic and optical products		
27	Manufacture of electrical equipment	C	
28	Manufacture of machinery and equipment n.e.c.		
29	Manufacture of motor vehicles, trailers and semitrailers	C	
30	Manufacture of other transport equipment		
31	Manufacture of furniture		
32	Other manufacturing	C	
33	Repair and installation of machinery and equipment		
35	Electricity, gas, steam and air conditioning supply	D	ELECTRICITY[...]
36	Water collection, treatment and supply	E	WATER [...] AND REMEDIATION ACTIVITIES
37	Sewerage	E	
38	Waste collection, treatment and disposal activities; materials recovery	E	

39	Remediation activities and other waste management services	E	
41	Construction of buildings	F	CONSTRUCTION
42	Civil engineering	F	
43	Specialised construction activities	F	
55	Accommodation	I	ACCOMMODATION[...]
56	Food and beverage service activities	I	
58	Publishing activities	J	INFORMATION AND COMMUNICATION
59	Motion picture, video and television programme production, sound recording and music publishing activities	J	
60	Programming and broadcasting activities	J	
61	Telecommunications	J	
62	Computer programming, consultancy and related activities	J	
63	Information service activities	J	
69	Legal and accounting activities	M	PROFESSIONAL, SCIENTIFIC AND TECHNICAL ACTIVITIES
70	Activities of head offices; management consultancy activities	M	
71	Architectural and engineering activities; technical testing and analysis	M	
72	Scientific research and development	M	
73	Advertising and market research	M	
74	Other professional, scientific and technical activities	M	
75	Veterinary activities	M	
77	Rental and leasing activities	N	ADMINISTRATIVE AND SUPPORT SERVICE ACTIVITIES
78	Employment activities	N	
79	Travel agency, tour operator reservation service and related activities	N	
80	Security and investigation activities	N	
81	Services to buildings and landscape activities	N	
82	Office administrative, office support and other business support activities	N	

3.B Country specific notes

Belgium

- Coverage: universe of firms that need to provide annual accounts
- Source: National Central Bank data. Information for small firms is completed using the VAT declaration which is compulsory for all firms.
- Unit of analysis: legal business entities, including affiliates of foreign businesses.
- Deflators: NACE 2 digit value added deflators provided by the National Bank of Belgium (Belgostat) and Eurostat.
- Threshold: none

- Additional filters: in the dataset that will be released, all cells for which less than 4 firms were used to make the computation will be removed.

Spain

- Source: Sample of mercantile firms.
- Coverage: Spanish Central Balance Sheet Database (CBSD) for most of large firms in Spain; Firm Registries for small and medium firms.
- Unit of analysis: mercantile firms (Sociedades Mercantiles) are included. Self-employed persons are excluded.
- Deflators: based on National Accounts data provided by National Statistical Office for 2000-2011; Banco de Espana estimates for 1995- 1999.
- Additional filters: due to confidentiality constraints, an observation was dropped if at least one of the following was fulfilled: (i) ≤ 3 firms were included or one firm accounted for more than 80% of VA.

3.C Common sector weight computation

The aggregation of sector data at the country level is done using a common set of sector weights for both countries. The reason is that there might be some small sectors which have outliers. Additionally, some countries have an over-representation of manufacturing firms. Country-specific weights might be highly affected by those issues. To partially mitigate these potential sources of bias we generate a set of common sector weights, computed as the share of value added each sector has in the whole economy (defined as the sum of all countries' values) in a specific year. More concretely: Let s denote the sector, i the country and t the year. The sector weight is defined as:

$$\theta_{st} = \sum_i VA_{sit} / \sum_s \sum_i VA_{sit} \quad (3.4)$$

In this case the weight does not depend on the specific country and it is allowed to differ across time. Hence given a certain variable ω_{sit} , defined for country i , sector s and time t , we would compute the corresponding country aggregate y_{it} (or across-sectors average) as follows:

$$y_{it} = \sum_s \theta_{st} \omega_{sit} \quad (3.5)$$

where θ_{st} is the sector weight.

3.D Additional material



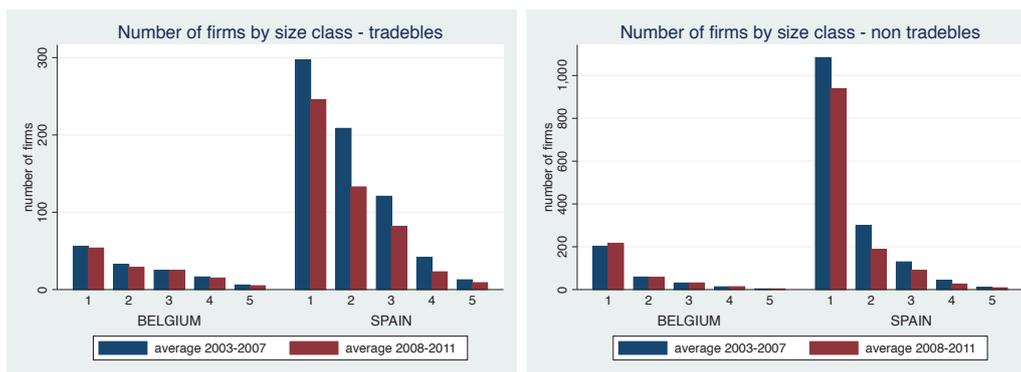


Table 3.D.1: Effect of the crisis on the OP gap in tradable/non-tradable sectors

	<i>Dependent variable: OP gap</i>	
	Belgium	Spain
$_Icrisis_1$	-0.0130 (-0.56)	0.00665 (0.22)
tradable	0.147*** (8.20)	0.150*** (6.33)
$_IcriXtrade_1$	0.0138 (0.39)	0.0133 (0.27)
N	857	986
F-stat	31.50	18.33
R-sq	0.0997	0.0530
Adj. R-sq	0.0966	0.0501

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.D.2: OP-gap

year	Belgium		Spain	
	tradable	non-tradable	tradable	non-tradable
1996	0.10	0.02	0.23	0.20
1997	0.04	0.01	0.23	0.19
1998	0.06	0.00	0.21	0.19
1999	0.06	0.03	0.20	0.17
2000	0.07	0.00	0.19	0.14
2001	0.06	-.01	0.21	0.04
2002	0.05	-0.01	0.20	0.04
2003	0.03	0.00	0.20	0.04
2004	0.04	-0.01	0.20	0.03
2005	0.05	-0.03	0.20	0.03
2006	0.05	-0.02	0.20	0.03
2007	0.06	-0.02	0.20	0.04
2008	0.07	-0.02	0.20	0.04
2009	0.03	0.03	0.14	0.13
2010	0.06	0.00	0.15	0.11
2011	0.05	0.06	0.16	0.14

Table 3.D.3: Average Firm Size, Manufacturing sectors

sector	1995-2007		2008-2011	
	Belgium	Spain	Belgium	Spain
10	24.99	25.59	24.28	25.56
11	96.89	28.30	80.96	24.84
12	100.52	909.95	68.19	251.90
13	41.68	18.67	35.89	11.74
14	20.74	18.33	16.59	16.08
15	27.14	15.01	25.69	12.53
16	16.56	12.76	17.27	10.74
17	59.66	36.47	59.58	35.79
18	13.63	11.29	13.17	9.44
19	432.97	635.23	529.20	633.50
20	114.03	40.08	102.98	35.69
21	215.19	126.54	264.47	153.91
22	44.34	31.73	44.45	29.58
23	34.66	27.17	31.70	22.14
24	217.25	57.04	210.89	48.23
25	20.28	14.43	19.81	12.10
26	78.90	40.74	68.68	22.62
27	63.02	54.69	61.73	52.11
28	42.36	26.30	43.83	22.10
29	236.76	181.35	201.02	174.44
30	116.30	224.40	112.35	160.93
31	19.31	14.39	15.73	10.70
32	10.96	12.59	10.75	10.40
33	31.80	11.28	29.90	9.43

Table 3.D.4: Average labor productivity, Manufacturing sectors

sector	1995-2007		2008-2011	
	Belgium	Spain	Belgium	Spain
10	58.8	33.8	68.8	31.3
11	96.2	54.1	105.5	47.7
12	117.7	56.3	150.8	34.9
13	55.4	28.8	67.5	29.8
14	38.4	22.9	60.3	26.2
15	39.8	23.3	56.7	25.3
16	54.4	29.3	69.0	28.5
17	57.2	38.8	83.8	39.8
18	65.2	34.8	74.9	33.8
19	173.0	92.5	268.3	77.5
20	96.1	53.9	97.6	45.3
21	82.7	62.5	84.2	72.8
22	60.3	39.3	91.5	38.8
23	67.9	39.9	64.4	34.4
24	77.7	49.9	87.7	41.9
25	60.6	37.4	70.3	34.3
26	64.1	35.4	67.4	40.8
27	58.0	44.2	60.7	41.1
28	60.7	43.2	64.2	40.8
29	48.3	39.6	47.8	39.2
30	47.9	42.3	84.2	43.5
31	54.9	26.1	52.4	24.8
32	59.6	31.3	57.7	31.8
33	73.8	33.7	50.9	34.1

Bibliography

Acemoglu, D., Akcigit, U., Bloom, N., and Kerr, W. R. (2013). Innovation, reallocation and growth. Technical report, National Bureau of Economic Research.

Akerberg, D., Caves, K., Aghion, P., Fally, T., Altomonte, C., Barba Navaretti, G., di Mauro, F., Altomonte, C., Aquilante, T., Altomonte, C., et al. (2012). Public policy and resource allocation: Evidence from firms in oecd countries. Technical report, OECD Economics Department Working Papers 996, OECD Publishing.

Albert, C., Garcia-Serrano, C., and Hernanz, V. (2005). Firm-provided training and temporary contracts. *Spanish Economic Review*, 7(1):67–88.

Altomonte, C. and Aquilante, T. (2012). The eu-efige/bruegel-unicredit dataset. Technical report, Bruegel working paper.

Altomonte, C., Barba Navaretti, G., Di Mauro, F., and Ottaviano, G. (2011). Assessing competitiveness: how firm-level data can help. Technical report, Bruegel Policy Contribution.

Andrews, D. and Cingano, F. (2012). Regulation, allocative efficiency and productivity: Evidence from firms in oecd countries. Technical report, OECD Economics Department Working Paper.

Angrist, J., Bettinger, E., and Kremer, M. (2006). Long-Term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia. *American Economic Review*, 96(3):847–862.

Arnold, J. M., Nicoletti, G., and Scarpetta, S. (2011). Does anti-competitive regulation matter for productivity? evidence from european firms.

- Bacow, L. S., Bowen, W. G., Guthrie, K. M., Lack, K. A., and Long, M. P. (2012). Barriers to adoption of online learning systems in us higher education. *New York, NY: Ithaka S+ R*.
- Banerjee, A. V. and Duflo, E. (2005). Growth theory through the lens of development economics. *Handbook of economic growth*, 1:473–552.
- Bardazzi, R. and Duranti, S. (2012). Atypical contracts and italian firms'labour productivity. Technical report, mimeo.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2004). Microeconomic evidence of creative destruction in industrial and developing countries.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2009). Measuring and analyzing cross-country differences in firm dynamics. In *Producer dynamics: New evidence from micro data*, pages 15–76. University of Chicago Press.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2013a). Cross-country differences in productivity: The role of allocation and selection. *The American Economic Review*, 103(1):305–334.
- Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2013b). Cross-country differences in productivity: The role of allocation and selection. *The American Economic Review*, 103(1):305–334.
- Bassanini, A., Nunziata, L., and Venn, D. (2009). Job protection legislation and productivity growth in oecd countries. *Economic Policy*, 24(58):349–402.
- Batten, D. F. (2000). Discovering artificial economics. *How agents learn and economies evolve*.
- Bentolila, S., Cahuc, P., Dolado, J. J., and Le Barbanchon, T. (2012). Two-tier labour markets in the great recession: France versus spain*. *The Economic Journal*, 122(562):F155–F187.
- Bergmann, B. R. (1973). Combining microsimulation and regression: A prepared regression of poverty incidence on unemployment and growth. *Econometrica: Journal of the Econometric Society*, pages 955–963.

- Bergmann, B. R. (1990). Micro-to-macro simulation: a primer with a labor market example. *The Journal of Economic Perspectives*, pages 99–116.
- Beyer, H.-G. and Schwefel, H.-P. (2002). Evolution strategies—a comprehensive introduction. *Natural computing*, 1(1):3–52.
- Blanchard, O. and Landier, A. (2002). The perverse effects of partial labour market reform: fixed-term contracts in france. *The Economic Journal*, 112(480):F214–F244.
- Boeri, T. (2011). Institutional reforms and dualism in european labor markets. *Handbook of labor economics*, 4:1173–1236.
- Boeri, T. and Garibaldi, P. (2007). Two tier reforms of employment protection: a honeymoon effect?*. *The Economic Journal*, 117(521):F357–F385.
- Boeri, T. and Garibaldi, P. (2010). Financial shocks and the labor markets: should economic policy save jobs?
- Bonk, C. J. and Graham, C. R. (2012). *The handbook of blended learning: Global perspectives, local designs*. John Wiley & Sons.
- Bosch, T. E. (2009). Using online social networking for teaching and learning: Facebook use at the university of cape town. *Communicatio: South African Journal for Communication Theory and Research*, 35(2):185–200.
- Boudreau, J. W. (2010). Stratification and growth in agent-based matching markets. *Journal of Economic Behavior & Organization*, 75(2):168–179.
- Bowen, W. G., Chingos, M. M., Lack, K. A., and Nygren, T. I. (2014). Interactive learning online at public universities: Evidence from a six-campus randomized trial. *Journal of Policy Analysis and Management*, 33(1):94–111.
- Bowen, W. G. and Ithaka, S. (2012). *Interactive learning online at public universities: Evidence from randomized trials*. Ithaka S+ R.
- Bradford, G. R. (2011). A relationship study of student satisfaction with learning online and cognitive load: Initial results. *The Internet and Higher Education*, 14(4):217–226.

- Brown, B. W. and Liedholm, C. E. (2002). Can Web Courses Replace the Classroom in Principles of Microeconomics? *American Economic Review*, 92(2):444–448.
- Brugiavini, A., Carraro, C., and Kovacic, M. (2014). Academic achievements: Grades versus duration. *Available at SSRN 2506602*.
- Cabrales, A., Dolado, J. J., and Mora, R. (2014). Dual labour markets and (lack of) on-the-job training: Piac evidence from Spain and other EU countries.
- Cahuc, P. and Carcillo, S. L. (2011). Is short-time work a good method to keep unemployment down?
- Cahuc, P., Charlot, O., and Malherbet, F. (2012). Explaining the spread of temporary jobs and its impact on labor turnover.
- Cahuc, P. and Postel-Vinay, F. (2002). Temporary jobs, employment protection and labor market performance. *Labour Economics*, 9(1):63–91.
- Cappellari, L., Dell’Ariaga, C., and Leonardi, M. (2012). Temporary employment, job flows and productivity: A tale of two reforms. *The Economic Journal*, 122(562):F188–F215.
- Cappelli, P. and Chauvin, K. (1991). An interplant test of the efficiency wage hypothesis. *The Quarterly Journal of Economics*, pages 769–787.
- Centeno, M. and Novo, Á. A. (2012). Excess worker turnover and fixed-term contracts: Causal evidence in a two-tier system. *Labour Economics*, 19(3):320–328.
- Chen, P.-S. D., Lambert, A. D., and Guidry, K. R. (2010). Engaging online learners: The impact of web-based learning technology on college student engagement. *Computers and Education*, 54(4):1222–1232.
- CompNet, T. F. (2014a). Micro-based evidence of EU competitiveness: The compnet database. Technical report.
- CompNet, T. F. (2014b). Micro-based evidence of EU competitiveness: The compnet database. Technical report.

- Damiani, M., Pompei, F., and Ricci, A. (2014). Temporary job protection and productivity growth in eu economies. *International Labour Review*.
- Daveri, F. and Parisi, M. L. (2010). Experience, innovation and productivity.
- Dawid, H., Gemkow, S., Harting, P., Kabus, K., Wersching, K., and Neugart, M. (2008). Skills, innovation, and growth: an agent-based policy analysis. *Jahrbücher für Nationalökonomie und Statistik*, pages 251–275.
- Dawid, H., Gemkow, S., Harting, P., and Neugart, M. (2009). On the effects of skill upgrading in the presence of spatial labor market frictions: an agent-based analysis of spatial policy design. *Journal of Artificial Societies and Social Simulation*, 12(4):5.
- Dawid, H., Gemkow, S., Harting, P., and Neugart, M. (2012). Labor market integration policies and the convergence of regions: the role of skills and technology diffusion. *Journal of Evolutionary Economics*, 22(3):543–562.
- Dawid, H., Harting, P., and Neugart, M. (2014). Economic convergence: policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control*, 44:54–80.
- Dawid, H. and Neugart, M. (2011). Agent-based models for economic policy design. *Eastern Economic Journal*, 37(1):44–50.
- Deming, D. J., Goldin, C., Katz, L. F., and Yuchtman, N. (2015). Can online learning bend the higher education cost curve? Technical report, National Bureau of Economic Research.
- Dew-Becker, I. and Gordon, R. J. (2008). The role of labor market changes in the slowdown of european productivity growth. Technical report, National Bureau of Economic Research.
- Di Pace, F., Mitra, K., and Zhang, S. (2014). Adaptive learning and labour market dynamics. Technical report, Centre for Dynamic Macroeconomic Analysis.
- Dolado, J. J., Jansen, M., and Jimeno-Serrano, J. F. (2005). Dual employment protection legislation: a framework for analysis.

- Dolado, J. J., Ortigueira, S., and Stucchi, R. (2012). Does dual employment protection affect tfp? evidence from spanish manufacturing firms.
- Dziuban, C. and Moskal, P. (2011). A course is a course is a course: Factor invariance in student evaluation of online, blended and face-to-face learning environments. *The Internet and Higher Education*, 14(4):236–241.
- Eliasson, G. (1977). Competition and market processes in a simulation model of the swedish economy. *The American Economic Review*, pages 277–281.
- Engellandt, A. and Riphahn, R. T. (2005). Temporary contracts and employee effort. *Labour economics*, 12(3):281–299.
- Fagiolo, G., Dosi, G., and Gabriele, R. (2004). Matching, bargaining, and wage setting in an evolutionary model of labor market and output dynamics. *Advances in Complex Systems*, 7(02):157–186.
- Fagiolo, G., Dosi, G., and Gabriele, R. (2005). Towards an evolutionary interpretation of aggregate labor market regularities. In *Entrepreneurships, The New Economy and Public Policy*, pages 223–252. Springer.
- Faini, R. and Sapir, A. (2005). Un modello obsoleto? crescita e specializzazione dell'economia italiana. *Oltre il declino, Bologna: il Mulino*, pages 19–77.
- Figlio, D., Rush, M., and Yin, L. (2013). Is It Live or Is It Internet?; Experimental Estimates of the Effects of Online Instruction on Student Learning. *Journal of Labor Economics*, 31(4):763 – 784.
- Figlio, D. N., Rush, M., and Yin, L. (2010). Is it Live or is it Internet? Experimental Estimates of the Effects of Online Instruction on Student Learning. NBER Working Papers 16089, National Bureau of Economic Research, Inc.
- Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425.
- Freeman, R. B. (1998). War of the models: which labour market institutions for the 21st century? *Labour Economics*, 5(1):1–24.

- Gabriele, R. (2002a). Labor market dynamics and institutions: an evolutionary approach. Technical report, LEM Working Paper Series.
- Gabriele, R. (2002b). Labor market dynamics and institutions: an evolutionary approach. Technical report, LEM Working Paper Series.
- Gagliarducci, S. (2005). The dynamics of repeated temporary jobs. *Labour Economics*, 12(4):429–448.
- Gemkow, S. and Neugart, M. (2011). Referral hiring, endogenous social networks, and inequality: an agent-based analysis. *Journal of Evolutionary Economics*, 21(4):703–719.
- Ghignoni, E. (2009). Temporary contracts, employees' effort and labour productivity: the evidence for italy. *Economia politica*, 26(2):279–314.
- Guadalupe, M. (2003). The hidden costs of fixed term contracts: the impact on work accidents. *Labour economics*, 10(3):339–357.
- Güell, M. and Petrongolo, B. (2007). How binding are legal limits? transitions from temporary to permanent work in spain. *Labour Economics*, 14(2):153–183.
- Haber, J. (2014). *MOOCs*. MIT Press.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pages 153–161.
- Hernández-Julián, R. and Peters, C. (2012). Targeting teaching: Does the medium matter? online versus paper coursework. *Southern Economic Journal*, 78(4):1333–1345.
- Hew, K. F. (2011). Students' and teachers' use of facebook. *Computers in Human Behavior*, 27(2):662–676.
- Hirsch, B. and Mueller, S. (2012). The productivity effect of temporary agency work: Evidence from german panel data*. *The Economic Journal*, 122(562):F216–F235.
- Ichino, A. and Riphahn, R. T. (2005). The effect of employment protection on worker effort: Absenteeism during and after probation. *Journal of the European Economic Association*, 3(1):120–143.

- Ioannides, Y. M. and Loury, L. D. (2004). Job information networks, neighborhood effects, and inequality. *Journal of economic literature*, pages 1056–1093.
- Joyce, T. J., Crockett, S., Jaeger, D. A., Altindag, O., and O’Connell, S. D. (2014). Does Classroom Time Matter? A Randomized Field Experiment of Hybrid and Traditional Lecture Formats in Economics. NBER Working Papers 20006, National Bureau of Economic Research, Inc.
- Junco, R. (2012). The relationship between frequency of facebook use, participation in facebook activities, and student engagement. *Computers & Education*, 58(1):162–171.
- Kabilan, M. K., Ahmad, N., and Abidin, M. J. Z. (2010). Facebook: An online environment for learning of english in institutions of higher education? *The Internet and Higher Education*, 13(4):179–187.
- Kirschner, P. A. and Karpinski, A. C. (2010). Facebook® and academic performance. *Computers in human behavior*, 26(6):1237–1245.
- Kleiber, C. and Zeileis, A. (2008). *Applied econometrics with R*. Springer.
- Kleinknecht, A., Oostendorp, R. M., Pradhan, M. P., and Naastepad, C. (2006). Flexible labour, firm performance and the dutch job creation miracle. *International Review of Applied Economics*, 20(02):171–187.
- Kwak, D. W., Menezes, F., and Sherwood, C. (2013). Assessing the impact of blended learning on student performance. Discussion Papers Series 494, School of Economics, University of Queensland, Australia.
- Kwak, D. W., Menezes, F. M., and Sherwood, C. (2014). Assessing the impact of blended learning on student performance. *Economic Record*.
- Leombruni, R. and Richiardi, M. (2005). Why are economists sceptical about agent-based simulations? *Physica A: Statistical Mechanics and its Applications*, 355(1):103–109.
- Lewkowicz, Z. and Kant, J.-D. (2008). A multiagent simulation of a stylized french labor market: Emergences at the micro level. *Advances in Complex Systems*, 11(02):217–230.
- Lotti, F. and Viviano, E. (2012). Why hiring temporary workers. *Rome: Banca d’Italia*.

- Mancini, M. (2011). Labor effort and temporary jobs: evidence for italy. Technical report, Università di Perugia, Dipartimento Economia, Finanza e Statistica.
- Martin, C. W. and Neugart, M. (2009). Shocks and endogenous institutions: an agent-based model of labor market performance in turbulent times. *Computational Economics*, 33(1):31–46.
- Means, B., Toyama, Y., Murphy, R., Bakia, M., and Jones, K. (2009). Evaluation of evidence-based practices in online learning: A meta-analysis and review of online learning studies. *US Department of Education*.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.
- Michie, J. and Sheehan, M. (1999). Hrm practices, r&d expenditure and innovative investment: evidence from the uk's 1990 workplace industrial relations survey (wirs). *Industrial and Corporate Change*, 8(2):211–234.
- Miller, J. H., Page, S. E., and LeBaron, B. (2008). *Complex adaptive systems: an introduction to computational models of social life*. Princeton University Press Princeton:.
- Neugart, M. (2004). Endogenous matching functions: An agent-based computational approach. *Advances in Complex Systems*, 7(02):187–201.
- Neugart, M. (2008). Labor market policy evaluation with ace. *Journal of Economic Behavior & Organization*, 67(2):418–430.
- Neugart, M., Richiardi, M., et al. (2012a). Agent-based models of the labor market. *LABORatorio R. Revelli working papers series*, 125.
- Neugart, M., Richiardi, M., et al. (2012b). Agent-based models of the labor market. *LABORatorio R. Revelli working papers series*, 125.
- Olley, G. S. and Pakes, A. (1992a). The dynamics of productivity in the telecommunications equipment industry. Technical report, National Bureau of Economic Research.
- Olley, G. S. and Pakes, A. (1992b). The dynamics of productivity in the telecommunications equipment industry. Technical report, National Bureau of Economic Research.

- Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.
- Pasek, J., Hargittai, E., et al. (2009). Facebook and academic performance: Reconciling a media sensation with data. *First Monday*, 14(5).
- Pellizzari, P. (2012). Facebook as an academic learning platform: A case study in Mathematics. Technical report, Department of Economics, University of Venice Ca, Foscari.
- Pérez, J. I. G., de Olavide, U. P., and Osuna, V. (2012). The effects of introducing a single open-ended contract in the spanish labour market. In *FEDEA/Fondazione Rodolfo Debenedetti/IZA Conference: Dual Labor Markets and the Single Contract*.
- Petrin, A. and Levinsohn, J. (2012). Measuring aggregate productivity growth using plant-level data. *The RAND Journal of Economics*, 43(4):705–725.
- Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic literature*, pages 390–431.
- Restuccia, D. and Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720.
- Richiardi, M. (2004). A search model of unemployment and firm dynamics. *Advances in Complex Systems*, 7(02):203–221.
- Richiardi, M. (2006). Toward a non-equilibrium unemployment theory. *Computational Economics*, 27(1):135–160.
- Richiardi, M. G. (2012). Agent-based computational economics: a short introduction. *The Knowledge Engineering Review*, 27(02):137–149.
- Riphahn, R. T. and Thalmaier, A. (2001). Behavioral effects of probation periods: An analysis of worker absenteeism/anreizeffekte der probezeit: Eine untersuchung von fehlzeiten bei arbeitnehmern. *Jahrbücher für Nationalökonomie und Statistik*, pages 179–201.
- Rogerson, R. (1997). Theory ahead of language in the economics of unemployment. *The Journal of Economic Perspectives*, pages 73–92.

- Ronchetti, M. (2010). Using video lectures to make teaching more interactive. *International Journal of Emerging Technologies in Learning (iJET)*, 5(2).
- Roth, A. E. and Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, 8(1):164–212.
- Shapiro, C. and Stiglitz, J. E. (1984). Equilibrium unemployment as a worker discipline device. *The American Economic Review*, pages 433–444.
- Silva, S. T., Valente, J. M., and Teixeira, A. A. (2012). An evolutionary model of industry dynamics and firms’ institutional behavior with job search, bargaining and matching. *Journal of Economic Interaction and Coordination*, 7(1):23–61.
- Tassier, T. and Menczer, F. (2001). Emerging small-world referral networks in evolutionary labor markets. *Evolutionary Computation, IEEE Transactions on*, 5(5):482–492.
- Tassier, T. and Menczer, F. (2008). Social network structure, segregation, and equality in a labor market with referral hiring. *Journal of Economic Behavior & Organization*, 66(3):514–528.
- Tealdi, C. (2011). Typical and atypical employment contracts: the case of italy.
- Tesfatsion, L. (2001a). Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control*, 25(3):281–293.
- Tesfatsion, L. (2001b). Structure, behavior, and market power in an evolutionary labor market with adaptive search. *Journal of Economic Dynamics and Control*, 25(3):419–457.
- Tesfatsion, L. (2002). *Hysteresis in an evolutionary labor market with adaptive search*. Springer.
- Toharia, L. and Serrano, J. F. J. (1996). Effort, absenteeism and fixed term employment contracts. *Revista Española de Economía*, 13(1):105–119.
- Tsallis, C. (2009). Learning with boltzmann–gibbs statistical mechanics. In *Introduction to Nonextensive Statistical Mechanics*, pages 19–36. Springer.

Xu, D. and Jaggars, S. S. (2013). The impact of online learning on students' course outcomes: Evidence from a large community and technical college system. *Economics of Education Review*, 37:46–57.

Acknowledgements

Carlo Fano e Laurie Pearlman, grazie per avermi educato, incoraggiato e lasciato poi a me la libertà di ogni scelta.

Ringrazio il Prof. Tito Boeri per avermi suggerito l'idea di studiare il tema del dualismo nel mercato del lavoro; grazie per poche parole ma sempre al posto giusto e per l'esempio affabile e determinato nel modo di interpretare il ruolo dell'economista.

Ringrazio il Prof. Paolo Pellizzari per avermi fatto pensare che la laurea potesse essere, oltre ad un traguardo, anche un punto di partenza. Grazie per essere stato un riferimento sempre presente in questo lungo percorso. Il suo ottimismo ed entusiasmo mi trasmettono la sua costante voglia di imparare cose nuove.

Ringrazio i Prof. Tommaso Nannicini, Nicola Pavoni, Alberto Alesina e Thomas Le Baranchon per i loro preziosi commenti e suggerimenti.

Carlo Maria Gallimberti e Volker Lindenthal, oltre al resto, hanno reso via Bellarmino davvero una seconda casa, anche nelle lunghe notti di lavoro su paper e problemset; non sarebbe stato lo stesso senza la loro amicizia.

So many books, so little time.

Frank Zappa